Operational Aspects of C/C++ Concurrency

July 12, 2016

Anton Podkopaev
Saint Petersburg State University and JetBrains Inc., Russia
a.podkopaev@2009.spbu.ru

Ilya Sergey
University College London, UK
i.sergey@ucl.ac.uk

Aleksandar Nanevski
IMDEA Software Institute, Spain
aleks.nanevski@imdea.org

arXiv:1606.01400v2 [cs.PL] 9 Jul 2016

Abstract

In this work, we present a family of operational semantics that gradually approximates the realistic program behaviors in the C/C++11 memory model. Each semantics in our framework is built by elaborating and combining two simple ingredients: *viewfronts* and *operation buffers*. Viewfronts allow us to express the *spatial* aspect of thread interaction, *i.e.*, which values a thread can read, while operation buffers enable manipulation with the *temporal* execution aspect, *i.e.*, determining the order in which the results of certain operations can be observed by concurrently running threads.

Starting from a simple abstract state machine, through a series of gradual refinements of the abstract state, we capture such language aspects and synchronization primitives as *release/acquire* atomics, *sequentially-consistent* and *non-atomic* memory accesses, also providing a semantics for *relaxed* atomics, while avoiding the Out-of-Thin-Air problem. To the best of our knowledge, this is the first formal and *executable* operational semantics of C11 capable of expressing all essential concurrent aspects of the standard.

We illustrate our approach via a number of characteristic examples, relating the observed behaviors to those of standard litmus test programs from the literature. We provide an executable implementation of the semantics in PLT Redex, along with a number of implemented litmus tests and examples, and showcase our prototype on a large case study: randomized testing and debugging of a realistic Read-Copy-Update data structure.

1. Introduction

Memory models describe the behavior of multithreaded programs, which might concurrently access shared memory locations. The best studied memory model is *sequential consistency* (SC) [28], which assumes a total order on all memory accesses (*i.e.*, read and write operations) in a single run of a concurrent program, therefore, ensuring that the result of each read from a location is a value that was stored by the last preceding write to the very same location.

However, sequential consistency falls short when describing the phenomena, observed in concurrent programs running on modern processor architectures, such x86, ARM, and PowerPC, and resulting from store buffering [50] and CPU- and compiler-level optimizations, *e.g.*, rearranging independent reads and writes [21]. *Relaxed memory models* aim to capture the semantics of such programs and provide suitable abstractions for the developers to reason about their code, written in a higher-level language, independently from the hardware architecture it is going to be executed on.

The most prominent example of a relaxed memory model is the C11 model, introduced by the C/C++ 2011 standards [2, 3] and describing the behavior of concurrent C/C++ programs. It defines a number of memory accesses, implementing different synchronization policies and having corresponding performance costs. For instance, *SC-atomics* provide the SC-style total ordering between reads and writes to the corresponding memory locations, while *release/acquire* (RA) accesses implement only partial one-way synchronization, but are cheaper to implement. Finally, *relaxed* accesses are the cheapest in terms of performance, but provide the weakest synchronization guarantees.

Existing formalizations of the *full* C11 memory model adopt an axiomatic style, representing programs by sets of *consistent executions* [5–7]. Each execution can be thought of as a graph, whose nodes are read/write-accesses to memory locations. The edges of the graph represent various orders between operations (*e.g.*, total orders between SC-atomics and operations in a single thread), some of which might be partial. Defined this way, the executions help one to answer questions of the following kind: "*Can the value X be read from the memory location L at the point R of the program P?*"

This axiomatic whole-program representation makes it difficult to think of C11 programs in terms of step-by-step executions of a program on some abstract machine, making it non-trivial to employ these semantic approaches for the purposes of testing, debugging and compositional symbolic reasoning about programs, *e.g.*, by means of type systems and program logics. Recently, several attempts have been made to provide a more operational semantics for C/C++ concurrency, however, all the approaches existing to date focus on a specific subset of C11, *e.g.*, release/acquire/SC synchronization [26,45] or relaxed atomics [38], without providing a uniform framework accommodating all features of the standard.

In this work, we make a step towards providing a simple, yet uniform foundations for accommodating all of the essential aspects of the C11 concurrency, and describe a framework for defining operational semantics capturing the expected behaviors of concurrent executions observed in realistic C/C++ programs, while prohibiting unwelcome outcomes, such as Thin-Air executions. The paramount idea of our constructions is maintaining a *rich program state*, which is a subject of manipulation by concurrent threads, and is represented by a combination of the following two ingredients.

Ingredient 1: Viewfronts for threads synchronization We observe that, assuming a total ordering of writes to each particular shared memory location, we can consider a state to be a collection of *per-location histories*, representing totally-ordered updates—an idea adopted from the recent works on logics for SC concurrency [41]. We introduce the notion of *viewfronts* as a way to account for the phenomenon of particular threads having specific, yet consistent, views to the global history of each shared location, similarly to the way vector clocks are used for synchronization in distributed systems [30]. We then consider various flavors of C11 atomicity as ways to "partially align" viewfronts of several threads.

Ingredient 2: Operation buffers for speculative executions The mechanism of relaxed atomic accesses in C11 allows for speculative reordering or removing of operations, involving them, in particular threads. In order to formally define the resulting temporal phenomena, observed by concurrently running threads (which can see some values appearing "out-of-order"), we need to capture a speculative nature of such computations. As an additional challenge, the semantics has to prohibit so-called Out-of-Thin-Air executions, in which results appear out of nowhere. We solve both problems by adopting the notion of *operation buffers* from earlier works on relaxed memory models [10, 11, 16], and enhancing it with *nesting* structure as a way to account for conditional speculations.

While simple conceptually, the two described ingredients, when combined, allow us to capture precisely the behavior of standard C11 synchronization primitives (including *consume*-reads), desired semantics of relaxed atomics, as well as multiple aspects of their interaction, by elaborating the treatment of viewfronts and buffers.

The C11 standard is intentionally designed to be very general and allow for multiple behaviors. However, particular compilation schemes into different target architectures might focus only on specific subsets of the enumerated features. To account for this diversity, our framework comes in an *aspect-oriented* flavor: it allows one to "switch on and off" specific aspects of C11 standard and to deal only with particular sets of allowed concurrent behaviors.

1.1 Contributions and outline

We start by outlining the basic intuition and illustrating a way of handling C11's RA-synchronization and speculative executions, introducing the idea of thread-specific viewfronts and operation buffers in Section 2. Section 3 demonstrates more advanced aspects of C11 concurrency expressed in our framework. Section 4 gives a formal definition of the operational model for C11, which is our central theoretical contribution. Section 5 describes evaluation of our semantics implemented in the PLT Redex framework [17, 25]. We argue for the adequacy of our constructions with respect to the actual aspects of C11 using a large corpus of *litmus test* programs, adopted from the earlier works on formalizing C11 concurrency. To do so, we summarize the described operational aspects of concurrent program behavior in C11, relating them to outputs of litmus tests. In Section 6, we showcase our operational model by tackling a large realistic example: testing and debugging several instances of a concurrently used Read-Copy-Update data structure [31, 34], implemented under relaxed memory assumptions. Our approach successfully detects bugs in the cases when the employed synchronization primitives are not sufficient to enforce the atomicity requirements, providing an execution trace, allowing the programmer to reproduce the problem. We compare to the related approaches to formalizing operational semantics for relaxed memory in general and for C11 in particular in Section 7, and conclude with a discussion of the future work in Section 8.

2. Overview and Intuition

We start by building the intuition for the program behaviors one can observe in the C11 relaxed memory model.

The code below implements the *message passing* pattern, where one of the two parallel threads waits for the notification from another one, and upon receiving it proceeds further with execution.

```
[f] := 0; [d] := 0;

[d] := 5;  ||  repeat [f] end;
[f] := 1;  ||  r = [d]
```

The identifiers in square parentheses (*e.g.*, `[f]`) denote accesses (*i.e.*, writes and reads) to *shared mutable* memory locations, subject to concurrent manipulation, whereas plain identifiers (*e.g.*, `r`) stand for *thread-local* variables. In a sequentially consistent setting, assuming that reads and writes to shared locations happen atomically, the right thread will not reach the last assignment to `r` until the left thread sets the flag `f` to be 1. This corresponds to the "message passing" idiom, and, hence, by the moment `[f]` becomes 1, `d` will be pointing to 5. so by the end of the execution, `r` must be 5.

$$
\begin{array}{l}
[\mathtt{f}]_{na} := 0;\ [\mathtt{d}]_{na} := 0; \\
[\mathtt{d}]_{na} := 5; \ \Big\| \ \mathtt{repeat}\ [\mathtt{f}]_{acq}\ \mathtt{end}; \\
[\mathtt{f}]_{rel} := 1; \ \Big\| \ \mathtt{r} = [\mathtt{d}]_{na}
\end{array}
$$

Figure 1. Release/acquire message passing (MP_rel+acq+na).

In a more realistic setting of C/C++ concurrent programming, it is not sufficient to declare all accesses to `[f]` and `[d]` as atomic: depending on particular *ordering annotations* on read/write accesses (*e.g.*, *relaxed*, *SC*, *release/acquire etc*) the outcome of the program might be different and, in fact, contradictory to the "natural" expectations. For instance, annotating all reads and writes to `[f]` and `[d]` as *relaxed* might lead to `r` being 0 at the end, due to the compiler and CPU-level optimizations, rearranging instructions of the left thread with no explicit dataflow dependency or, alternatively, assigning the value of `[d]` to `r` in the right thread speculatively.

One way to avoid these spurious results is to enforce stronger synchronization guarantees between specific reads and writes in a program using release/acquire order annotations. For instance, to ensure the "natural" behavior of the message-passing idiom, the program from above can be annotated as in Figure 1.

In the modified program, all accesses to the location `d` are now annotated as *non-atomic*, which means racy concurrent manipulations with them are considered run-time errors. What is more important, the write to `f` in the left thread is now annotated with `rel` modifier, which "publishes" the effects of the previous operations, making them *observable* by concurrently running threads after the assigned to `f` value 1 is read by them. Furthermore, the read from `f` in the right thread is now annotated with `acq`, preventing the operations following it in the same thread from taking effect before the read itself takes place. Together, the release/acquire modifiers in Figure 1 create a *synchronization order* we are seeking for and ensure that the second assignment to `r` will only take place after the `repeat`-loop terminates, hence the observed value of `f` is 1, which implies that the observed value of `d` is 5, thanks to the release-write, so the final value of `r` is also 5. Notice that there is also *no* race between the two concurrent non-atomic accesses to `d`, as those are clearly separated in time, thanks to the synchronization between release/acquire accesses to `f`.

Axiomatic semantics and execution orders The state-of-the-art formalization [7] of C11 defines semantics for program execution as a set of graphs, where nodes denote memory accesses (*i.e.*, reads and writes) with particular input/result values, and edges indicate ordering relations between them.[1]

One instance of an execution graph for MP_rel+acq+na is shown in Figure 2. The edges labelled by sb indicate a natural program order, reconstructed from the program's syntax. The green edge marked sw indicates the *synchronizes-with* relation, which arises dynamically between a release-write and acquire-read of the same value from the same location. The transitive closure of the union of the sb and sw relations is called *happens-before* relation (hb) and is central for defining the observed behaviors. In particular, a value X, written at a point R_1 of a program, can be only read at a point R_2 if there is no hb-ordering between the corresponding read event in R_2 and write event in R_1. That, the *read-from* order (rf) must not contradict the hb order.

[1] For illustrative purposes, here we employ a version of execution graphs [49] with additional explicit nodes for spawning and joining threads.

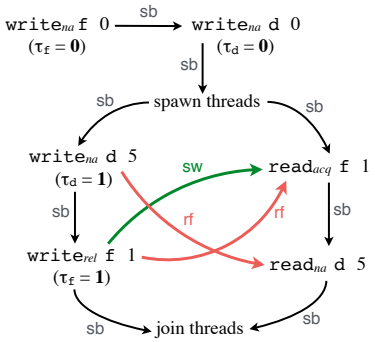

Figure 2. Execution orders in the message passing example.

The C11 standard defines a number of additional axioms, specifying the consistent execution graphs. In particular, all write actions on an atomic location (*i.e.*, such that it is not accessed non-atomically) must be totally ordered via *modification ordering* relation (mo), which is consistent with hb, restricted to this location: if a write operation W_2 to a location ℓ has a greater timestamp than W_1, also writing to ℓ, then W_1 cannot be hb-ordered *after* W_2. The mo relation is shown in Figure 2 via per-location timestamps τ_-, incremented as new values are stored. Moreover, a read event R_1 cannot read from W_1, if R_1 is hb-ordered after W_2, *i.e.*, it's aware of the later write W_2, as illustrated by rf-edges in Figure 2, preventing the case of reading of 0 from d.

2.1 Synchronizing threads' knowledge via viewfronts

The *read-from* relation determines the results of reads in specific program locations depending on preceding or concurrent writes by relying on the global sb and sw orderings, and restricting them according to the C11 axioms. To avoid the construction of global partial orders and provide an incremental operational formalism for executing concurrent C11 programs, we focus on the mo relation for specific locations, making it an inherent component of the program state. We call this state component a *history*: it contains totally ordered "logs" of updates to every shared memory location, indexed by timestamps (natural numbers). The history is objective (*i.e.*, *global*): it contains information about all updates to shared locations, as they took place during the execution.

However, the way threads "see" the history with respect to particular locations is subjective (*i.e.*, *thread-local*): each thread has its own knowledge of what is the "latest" written value to each location. Thus, the value a thread actually reads can be written *no earlier* than what the thread considers to be the location's latest value. To formalize this intuition, we define the notion of *viewfronts*.

A *viewfront* is a partial function from memory locations to natural numbers, representing timestamps in the corresponding location's part of the history. A thread's viewfront represents its knowledge of what were the timestamps of last written values to the relevant locations that it is aware of. When being a subject of a release-write, a location will store the viewfront of a writing thread, which we will refer to as *synchronization front*, in addition to the actual value being written. Symmetrically, another thread performing a synchronized load (*e.g.*, acquire-read) from the location will update its viewfront via the one "stored" in the location. Viewfronts, when used for expressing release/acquire synchronization, are reminiscent to *vector time frames* [39], but are used differently to express more advanced aspects of C11 atomicity (see Section 3 for details).

The following table represents the history and the threads' viewfronts for the example from Figure 1 at the moment the left thread has already written 1 and 5 to f and d, but the right thread has not yet exited the **repeat**-loop.

```
[x]_rlx := 0; [y]_rlx := 0;
r1 = [y]_rlx ; || r2 = [x]_rlx ;
[x]_rlx := 1   || [y]_rlx := 1
```

Figure 3. An example with early reads (LB_rlx).

τ	f		d	
0	0, $\bot$	R	0, $\bot$	R
1	1, (f $\mapsto$ **1**, d $\mapsto$ **1**)	L	5, $\bot$	L

The values in the first column (**0**, **1**) are *timestamps*, ascribing total order mo to the values written to a certain location (f and d correspondingly). The remaining columns capture the sequence of updates of the locations, each update represented as a pair of a value and a stored synchronization front. Viewfronts form a lattice, as partial maps from locations to timestamps, with $\bot = \emptyset$. The $\bot$ fronts are used for location updates corresponding to non-atomic stores. The front (f $\mapsto$ **1**, d $\mapsto$ **1**) was stored to f upon executing the corresponding release-write, capturing the actual viewfront L of the left thread. The right thread's viewfront R indicates that it still considers f and d to be at least at timestamp **0**, and, hence, can observe their values at timestamps larger or equal than **0**.

Let us explore a complete execution trace of the program from Figure 1. The initial state looks as follows:

τ	f	d
0	-	-

anf after the parent thread executes '$[\mathtt{f}]_{na} := 0$' and '$[\mathtt{d}]_{na} := 0$', it becomes:

τ	f		d	
0	0, $\bot$	P	0, $\bot$	P

Two subthreads are spawned, inheriting the parent's viewfront:

τ	f			d		
0	0, $\bot$	L	R	0, $\bot$	L	R

The left subthread performs '$[\mathtt{d}]_{na} := 5$', incrementing the timestamp τ of d and updating its own viewfront L:

τ	f			d	
0	0, $\bot$	L	R	0, $\bot$	R
1	-			5, $\bot$	L

Next, the left thread executes '$[\mathtt{f}]_{rel} := 1$', updating its viewfront and simultaneously storing it to the **1**-entry f:

τ	f		d	
0	0, $\bot$	R	0, $\bot$	R
1	1, (f $\mapsto$ **1**, d $\mapsto$ **1**)	L	5, $\bot$	L

The right thread can read the values of f, stored no later than its viewfront R indicates, thus, eventually it will perform the acquire-read from f with $\tau_f = \mathbf{1}$, updating its R correspondingly:

τ	f			d		
0	0, $\bot$			0, $\bot$		
1	1, (f $\mapsto$ **1**, d $\mapsto$ **1**)	L	R	5, $\bot$	L	R

Now the right thread's viewfront is updated with respect to the latest store to d, it reads 5 from it, and the threads join:

τ	f		d	
0	0, $\bot$		0, $\bot$	
1	1, (f $\mapsto$ **1**, d $\mapsto$ **1**)	P	5, $\bot$	P

2.2 Speculating with operation buffers

Relaxed atomics in C11 allow for speculative program optimizations, which might result in out-of-order behaviors, observed during concurrent executions under weak memory assumptions.

As a characteristic example of such a phenomenon, consider the program in Figure 3. The C11 standard [3], as well as its axiomatic

$[\mathtt{x}]_{rlx} := 0;\ [\mathtt{y}]_{rlx} := 0;\ [\mathtt{z}]_{rlx} := 0;$

if $[\mathtt{x}]_{rlx}$ **then** $[\mathtt{z}]_{rlx} := 1;$ $[\mathtt{y}]_{rlx} := 1$ **else** $[\mathtt{y}]_{rlx} := 1$ **fi**	**if** $[\mathtt{y}]_{rlx}$ **then** $[\mathtt{x}]_{rlx} := 1;$ **else** 0 **fi**

$\mathtt{res} := [\mathtt{z}]_{rlx}$

Figure 4. A program allowing if-speculations (SE_simple).

formal models [5,7,49], allow for the outcome $\mathtt{r1} = 1 \wedge \mathtt{r2} = 1$ by the end of its execution, as a result of rearranging instructions. Alas, our viewfront-based semantics cannot account for such a behavior: in order to be read from a location, a value should have been first stored into the history by some thread! However, in the example, it is either `x` or `y` that stores 1 (but not both) at the moments `r1` and `r2` were assigned. That is, while viewfront manipulation enables fine-grained control of *what* can be observed by threads, it does not provide enough flexibility to specify *when* effects of a particular thread's operations should become visible to concurrent threads.

To account for such anomalies of relaxed behaviors, we introduce *per-thread operation buffers*, which allow a thread to postpone an execution of an operation, "resolving" it later. An operation buffer itself is a queue of records, each of which contains an essential information for performing the corresponding postponed operation. For instance, for a postponed read action, a thread allocates a fresh *symbolic* value to substitute for a not-yet-resolved read result, and adds a tuple, containing the location and the symbolic value, to the buffer. For a write action the thread puts an another tuple, the location and the value to store to it, to the buffer. As it proceeds with the execution, the thread can non-deterministically resolve an operation from the buffer *if there is no operation before it*, which may affect its result, *e.g.*, a write to the same location, or an acquire-read changing the local viewfront. For instance, in Figure 3, buffering the effects of the two relaxed reads, $[y]_{rlx}$ and $[x]_{rlx}$, postpones their effects beyond the subsequent writes, enabling the desired outcome, as by the moment the reads are resolved, the corresponding 1's will be already stored to the history.

Nested buffers and speculative conditionals The idea of buffering operations for postponing their effects in a relaxed concurrency settings is not novel and has previously appeared in a number of related weak memory frameworks [10,11,13,16]. However, in our case it comes with a twist, making it particularly well suited for modelling C11 behaviors, while avoiding "bad" executions.

To illustrate this point, let us consider an example of a speculative optimization involving a conditional statements. Such optimizations are known to be difficult for modelling in relaxed concurrency [5,38]. For instance, in the program in Figure 4, the assignment $[y]_{rlx} := 1$ can be "pulled out" from both branches of the left thread's conditional statement, as it will be executed anyway, and, furthermore, it does not bear a data dependency with the possible preceding assignment $[y]_{rlx} := 1$. Such an optimization will, however, lead to interesting consequences: in the right thread, the conditional statement might succeed assigning 1 to `x`, therefore leading to the overall result $\mathtt{res} = 1$. This outcome relies on the fact that the optimization, which made $[y]_{rlx} := 1$ unconditionally visible to the right thread, was done speculatively, yet it has been justified later, since the same assignment would have been performed no matter which branch has been executed.

Luckily, to be able to express such a behavior, our buffer machinery requires only a small enhancement: *nesting*. In the semantics, upon reaching an if-then-else statement, whose condition's expression is a result of some preceding relaxed read, which is not yet resolved, we create a *tuple*, containing the symbolic representation of the condition as well as two empty buffers, to be filled

$[\mathtt{x}]_{rlx} := 0;\ [\mathtt{y}]_{rlx} := 0;$

if $[\mathtt{x}]_{rlx}$ **then** $[\mathtt{y}]_{rlx} := 1$ **else** 0 **fi**	**if** $[\mathtt{y}]_{rlx}$ **then** $[\mathtt{x}]_{rlx} := 1$ **else** 0 **fi**

$\mathtt{r1} = [\mathtt{x}]_{rlx};\ \mathtt{r2} = [\mathtt{y}]_{rlx}$

Figure 5. Program with C11-allowed Thin-Air behavior (OTA_if).

with postponed operations of the left and the right branches, correspondingly. The tuple is then added to the thread's main operation buffer.

More specifically, in the program SE_simple, the history after the three initial relaxed writes is as follows:

τ	x	y	z
0	0, ($\mathtt{x} \mapsto \mathbf{0}$)	0, ($\mathtt{y} \mapsto \mathbf{0}$)	0, ($\mathtt{z} \mapsto \mathbf{0}$)

The left thread then postpones reading from `x` and start the executing the **`if`** statement speculatively, with the following buffer:

$\langle\ \mathtt{a} = [\mathtt{x}]_{rlx};\ \mathbf{if}\ \mathtt{a}\ \langle\rangle\ \langle\rangle\ \rangle$

Proceeding to execute the two branches of the `if`-statement with focusing on the corresponding nested buffers, the left thread eventually fills them with the postponed commands:

$\langle\ \mathtt{a} = [\mathtt{x}]_{rlx};\ \mathbf{if}\ \mathtt{a}\ \langle [\mathtt{z}]_{rlx} := 1;\ [\mathtt{y}]_{rlx} := 1\rangle\ \langle [\mathtt{y}]_{rlx} := 1\rangle\ \rangle$

At this point the two sub-buffers contain the postponed write $[\mathtt{y}]_{rlx} := 1$, and no other postponed operations in the same buffers are in conflict with them. This allows the semantics to *promote* this write to the upper-level buffer (*i.e.*, the main buffer of the thread):

$\langle\ \mathtt{a} = [\mathtt{x}]_{rlx};\ [\mathtt{y}]_{rlx} := 1;\ \mathbf{if}\ \mathtt{a}\ \langle [\mathtt{z}]_{rlx} := 1\rangle\ \langle\rangle\ \rangle$

Next, the write is resolved, so its effect is visible to the right thread:

$\langle\ \mathtt{a} = [\mathtt{x}]_{rlx};\ \mathbf{if}\ \mathtt{a}\ \langle [\mathtt{z}]_{rlx} := 1\rangle\ \langle\rangle\ \rangle$

At that moment, the overall history looks as follows:

τ	x	y	z
0	0, ($\mathtt{x} \mapsto \mathbf{0}$)	0, ($\mathtt{y} \mapsto \mathbf{0}$)	0, ($\mathtt{z} \mapsto \mathbf{0}$)
1	-	1, ($\mathtt{y} \mapsto \mathbf{1}$)	-

Hence, the right thread can read 1 from the location `y`, take the **`then`** branch of the **`if`** statement, and perform the write to `x`:

τ	x	y	z
0	0, ($\mathtt{x} \mapsto \mathbf{0}$)	0, ($\mathtt{y} \mapsto \mathbf{0}$)	0, ($\mathtt{z} \mapsto \mathbf{0}$)
1	1, ($\mathtt{x} \mapsto \mathbf{1}$)	1, ($\mathtt{y} \mapsto \mathbf{1}$)	-

Now the left thread can resolve the postponed read $\mathtt{a} = [\mathtt{x}]_{rlx}$ obtaining 1 as its result and reducing the operation buffer:

$\langle\ \mathbf{if}\ 1\ \langle [\mathtt{z}]_{rlx} := 1\rangle\ \langle\rangle\ \rangle$

By evaluating the buffered **`if`** and resolving the write $[\mathtt{z}]_{rlx} := 1$:

τ	x	y	z
0	0, ($\mathtt{x} \mapsto \mathbf{0}$)	0, ($\mathtt{y} \mapsto \mathbf{0}$)	0, ($\mathtt{z} \mapsto \mathbf{0}$)
1	1, ($\mathtt{x} \mapsto \mathbf{1}$)	1, ($\mathtt{y} \mapsto \mathbf{1}$)	1, ($\mathtt{z} \mapsto \mathbf{1}$)

Reading from the latest record for `z` results in $\mathtt{res} = 1$.

The idea of nested buffers with promoting duplicating records from a lower to an upper level (under some dependency conditions) naturally scales for the case of nested if-statements.

On the Out-of-Thin-Air problem So what are the "bad" executions that should be prohibited by a meaningful semantics?

The C11 standard [2,3] and the axiomatic semantics [7] allow for so-called *Out-of-Thin-Air* (OTA) behaviors, witnessed by self-satisfying conditionals, such as the one represented by the program in Figure 5, which, according to the standard is allowed to end up with $\mathtt{r1} = \mathtt{r2} = 1$. Such behavior is, however, not observable on any of the major modern architectures (x86, ARM, and POWER), and considered as a flaw of the model [5,8], with researchers developing alternative semantics for relaxed atomics that avoid OTA [38].

$$[x]_{sc} := 0;\ [y]_{sc} := 0;$$
$$[x]_{sc} := 1; \quad \| \quad [y]_{sc} := 1;$$
$$r1 = [y]_{sc} \quad \| \quad r2 = [x]_{sc}$$

Figure 6. A program with SC synchronization (SB_sc).

Notice the only essential difference between the programs in Figure 4 and 5 is that in the former the write performed speculatively will *always* take place, whereas in the latter one the speculative writes in the then-branch might end up unjustified.

As we have previously demonstrated, our semantics supports the weak behavior of the program in Figure 4, and outlaws it for the program in Figure 5, as the conditions for promoting buffered operations in the if-branches will not be met in the latter case.

3. Advanced Aspects of C11 Concurrency

In this section, we elaborate and employ the ideas of viewfronts and operation buffers to adequately capture the remaining aspects of C11 concurrency. In particular, we *(i)* show how to extend the viewfront mechanism to support sequentially-consistent (SC) and non-atomic (NA) memory accesses, as well as consume-reads, (§§3.1–3.3); *(ii)* employ operation buffers to account for specific phenomena caused by sequentialization optimization (§3.4), and *(iii)* demonstrate the interplay between relaxed atomics and RA-synchronization (§3.5).

3.1 Sequentially-consistent memory accesses

To see the difference between SC-style and release/acquire-synchronization in C11, consider the program in Figure 6. All SC-operations are totally ordered with respect to each other, and the last of them is either read from x, or from y. Thus, the overall outcome $r1 = 0 \wedge r2 = 0$ is impossible. Replacing any of the *sc* modifiers by release or acquire in the corresponding writes and reads, makes $r1 = 0 \wedge r2 = 0$ a valid outcome, because there is no more the total order on all operations. In particular, the left subthread (with RA modifiers instead of SC) could still read 0 from y, as by its viewfront, which at that moment is $(x \mapsto \mathbf{1}, y \mapsto \mathbf{0})$.

In the axiomatic model [7,49], the restricted set of SC behaviors is captured by introducing an additional order sc and several axioms, requiring, in particular, consistency of sc with respect to hb and mo. In our operational setting, it means that an SC-read cannot read from a write with a *smaller* timestamp than the *greatest* timestamp of SC-writes to this location. To capture this requirement, we instrument the program state with an additional component—a *global* viewfront of sequentially consistent memory operations (σ_{sc}), which is being updated at each SC-write.

3.2 Non-atomic memory accesses and data races

Following the C11 standard [2, 3], our semantics does not draw a distinction between non-atomic and atomic locations (in contrast with the axiomatic model [7]). However, data races involving non-atomic memory *operations* (whose purpose is data manipulation, not thread synchronization) might result in an undefined behavior.

Consider the following two code fragments with data races on non-atomics. In the first case, a thread performs a na-read concurrently with a write to the same location.

$$[d]_{na} := 0;$$
$$[d]_{rlx} := 1 \quad \| \quad r = [d]_{na}$$

We can detect the data race, when the right subthread is executed after the left one, so it performs the na-read, while not being "aware" of the latest write to the same location. As our semantics constructs the whole state-space for all possible program executions, we will identify this data race on some execution path.

The second case is an opposite one: na-write and atomic read:

$$[p]_{na} := \mathtt{null};\ [d]_{na} := 0;\ [x]_{na} := 0;$$
$$[x]_{rlx} := 1; \quad \| \quad r1 = [p]_{con};$$
$$[d]_{na} := 1; \quad \| \quad \mathtt{if}\ \ r1 \mathrel{!}= \mathtt{null}$$
$$[p]_{rel} := d \quad \| \quad \mathtt{then}\ r2 = [r1]_{na};$$
$$\| \quad\quad r3 = [x]_{rlx}$$
$$\| \quad \mathtt{else}\ r2 = 0;\ r3 = 0$$
$$\| \quad \mathtt{fi}$$

Figure 7. An example with a consume read (MP_con+na_2).

$$[x]_{rlx} := 0;\ [y]_{rlx} := 0;$$
$$r1 = [y]_{rlx}; \quad \| \quad 0 \quad \| \quad r2 = [x]_{rlx}; \quad \| \quad 0$$
$$[z1]_{rlx} := r1 \quad \| \quad\quad\quad [z2]_{rlx} := r2$$
$$[x]_{rlx} := 1 \quad \| \quad\quad\quad [y]_{rlx} := 1$$

Figure 8. A program with non-flat thread joining (LB_rlx+join).

$$[d]_{na} := 0;$$
$$[d]_{na} := 1 \quad \| \quad r = [d]_{rlx}$$

It still has a data race involving a non-atomic access, which, however, we cannot detect by comparing threads' viewfronts. To identify data races of this kind, we extend the state with a global na-front, storing a timestamp of the last na-write to a location. Now, if the left thread executes its na-write first, the atomic read in the right one will not be aware of it, which will be manifested as a data race, thanks to the na-front.

3.3 Consume-reads

Unlike acquire-reads, consume-reads [33] do not update a thread's viewfront, but provide a synchronization front only for subsequent reads that are dereferencing their result.

Consider the code fragment in Figure 7. Here, we have message-passing of data stored in d via location p. The right thread employs a consume-read from p. In the case when it gets a non-**null** value (representing a pointer to d), it reads from it to r2, and after that from location x to r3. There might be three possible outcomes: $r1 = \mathtt{null} \wedge r2 = 0 \wedge r3 = 0$, $r1 = d \wedge r2 = 1 \wedge r3 = 1$, and $r1 = d \wedge r2 = 1 \wedge r3 = 0$. Changing the consume-read to an acquire one makes the last triple forbidden, as the right thread's viewfront would become up to date with both $[x]_{rlx} := 1$ and $[d]_{rlx} := 1$ after acquire-reading a non-**null** pointer value from p. At the same time, consume-read $r1 = [p]_{con}$ provides synchronization only for $r2 = [r1]_{rlx}$, which explicitly dereferences its result, but not for r3, which has no data-dependency with it.

Adding consume-reads to the semantics requires us to change the program syntax to allow *run-time annotations* on reads, which might be affected by consume ones. When a consume-read is executed, it retrieves some value/front-entry (v, σ) from the history, as any other read. Unlike an acquire-read, it does not update the thread's viewfront by the retrieved σ. Instead, it annotates all subsequent data-dependent reads by the front σ. Later, when these reads will be executed, they will join the front σ from the annotation with the thread's viewfront for computing the lower boundary on the relevant location's timestamp. The same process is applied to annotate data-dependent postponed reads in buffers, which might refer to the symbolic result of a consume-read.

3.4 Threads joining and synchronization

Once two threads join, it is natural to expect that all their postponed memory operations are resolved, *i.e.*, they have empty operation buffers. This is reflected in the axiomatic semantics [7] by an *additional-synchronizes-with* relation, which is a part of the *happens-before* relation. Thus, every memory action of joined threads *happens-before* actions which are syntactically after the

$[\mathtt{f}]_{na} := 0;\ [\mathtt{d}]_{na} := 0;\ [\mathtt{x}]_{na} := 0;$

$[\mathtt{d}]_{na} := 5;$ $[\mathtt{f}]_{rel} := 1;$ $[\mathtt{x}]_{rel} := 1;$ $[\mathtt{f}]_{rlx} := 2$	‖	**repeat** $[\mathtt{f}]_{acq} == 2$ **end**; $\mathtt{r1} := [\mathtt{d}]_{na};$ $\mathtt{r2} := [\mathtt{x}]_{rlx}$

Figure 9. Example of release sequence (MP_rel+acq+na+rlx_2).

$[\mathtt{x}]_{rlx} := 0;\ [\mathtt{y}]_{rlx} := 0;$

$\mathtt{r1} = [\mathtt{y}]_{rlx};$ $[\mathtt{x}]_{rel} := 1$	‖	$\mathtt{r2} = [\mathtt{x}]_{rlx};$ $[\mathtt{y}]_{rel} := 1$

Figure 10. Postponed relaxed reads and release-writes (LB_rel+rlx).

join point. Our semantics achieves this by merging viewfronts at join and forcing resolution of all operations in the buffers.

However, the C11 standard is intended to allow *sequentialization* optimization [48], *i.e.*, $S_1 \parallel S_2 \rightsquigarrow S_1 ; S_2$, making the previous assumption unsound. This is illustrated by the program in Figure 8, in which the parallel compositions with "idle" threads might be optimized by replacing them with non-idle parts. After such an optimization it is possible to observe 1s as result values of `z1` and `z2`. To account for this, we allow an alternative instance of a thread-joining policy, implemented as an aspect, which takes all possible interleavings of the threads' operation buffers (with an idle thread being a natural unit), thus achieving the behavior we seek for.

3.5 Relaxed atomics and synchronization

Interaction between relaxed atomics and RA-synchronization is particularly subtle, due to a number of ways they might affect the outcomes of each other. We identify these points of interaction and describe several design decisions, elaborating the structure of the state, so the requirements imposed by the C11 standard are met.

3.5.1 Relaxed writes and release sequences

Since a relaxed read cannot be used for synchronization, in our semantics it does not update the viewfront of the thread with a synchronization front from the history (as an acquire read does). However, when an acquire-read in thread T_2 reads a result of a relaxed write performed by thread T_1, it should be synchronized with a *preceding* release-write to the same location performed by thread T_1, if there is one. This observation follows the spirit of the axiomatic model [7], which defines a notion of *release sequence* between the writes of thread T_1.

For an example, let us take a look at Figure 9 presenting a modified version of the message-passing program. The only possible outcome for `r1` is 5, because when $[\mathtt{f}]_{acq}$ gets 2, it also becomes synchronized with $[\mathtt{f}]_{rel} := 1$, which precedes $[\mathtt{f}]_{rlx} := 2$ in the left thread. At the same time `r2` can be either 0, or 1: 0 is a possible outcome for `r2`, because $[\mathtt{f}]_{acq}$ synchronizes with $[\mathtt{f}]_{rel} := 1$, which precedes the $[\mathtt{x}]_{rel} := 1$ write, which therefore might be missed.

To express this synchronization pattern in our model, we instrument the state with per-thread *write-fronts*, containing information about last release-writes to locations performed by the thread. Upon a relaxed write, this information is used to retrieve a synchronization front from the history record with a timestamp equal to the write-front value, contributed by a preceding release-write.

3.5.2 Postponed relaxed operations and synchronization

Consider the program in Figure 10, which is similar to the example from Figure 3, but writes have release-modifiers. Since a release-write does not impose any restriction without a related acquire read, it is still possible to get the result `r1 = r2 = 1`. Therefore, our semantics allows to perform a release-write even if there are postponed reads from other locations in the thread's buffer.

$[\mathtt{x}]_{rlx} := 0;\ [\mathtt{y}]_{rlx} := 0;$

$\mathtt{r1} = [\mathtt{y}]_{acq};$ $[\mathtt{x}]_{rlx} := 1$	‖	$\mathtt{r2} = [\mathtt{x}]_{rlx};$ $[\mathtt{y}]_{rel} := 1$

Figure 11. Postponed relaxed reads and RA (LB_rel+acq+rlx).

$[\mathtt{x}]_{rlx} := 0;\ [\mathtt{y}]_{rlx} := 0;$

$[\mathtt{x}]_{rlx} := 1;$ $[\mathtt{y}]_{rel} := 2$	‖	$[\mathtt{y}]_{rlx} := 1;$ $[\mathtt{x}]_{rel} := 2$

$\mathtt{r1} = [\mathtt{x}]_{rlx};\ \mathtt{r2} = [\mathtt{y}]_{rlx}$

Figure 12. Postponed writes and release-writes (WR_rlx+rel).

The program in Figure 11 is more problematic, as it has release/acquire modifiers on accesses to `y`, and postponing the read $[\mathtt{x}]_{rlx}$ in the right thread might lead to `r1 = r2 = 1`. Notice that such an outcome is in conflict with the semantics of RA, as assigning 1 to `r1` would imply synchronization between $[\mathtt{y}]_{rel} = 1$ and $\mathtt{r1} = [\mathtt{y}]_{acq}$, thus the former happens *before* the latter one! Following the rules of RA-synchronization, $\mathtt{r2} = [\mathtt{x}]_{rlx}$ should happen before $[\mathtt{x}]_{rlx} = 1$ as well, making it impossible to read 1 into `r2`.

The problem is clear now: we need to prevent from happening the situations, when a read, postponed beyond a release-write W, is resolved after some concurrent acquire-read gets synchronized with W, as it might damage RA-synchronization sequences.

To achieve this, we instrument the state with a global list γ of triples that consist of: *(i)* a location ℓ, *(ii)* a timestamp τ of some executed write, and *(iii)* a symbolic value `x` of a read postponed beyond this write. When executing a release-write W, which stores a value to a location ℓ with a timestamp τ, for each thread-local postponed read, we globally record a triple $\langle \ell, \tau, \mathtt{x} \rangle$, where `x` is the symbolic value of the read. An acquire-read of the (ℓ, τ) history entry by another thread succeeds only if there are no $\langle \ell, \tau, \mathtt{x} \rangle$ left in γ for any symbolic value `x`. Resolving a postponed read with a symbolic value `x` removes all `x`-related entries from γ.

The program in Figure 12 (2+2W from [26]) demonstrates another subtlety, caused by interaction between postponed writes and RA-synchronization. According to the standard, `r1 = r2 = 1` is its valid outcome, and to achieve that in our semantics the corresponding relaxed writes should be committed to the history *after* the release ones. However, if the another (third) thread performs an acquire read from the history record of one of the release writes, it should become aware of the previous corresponding relaxed write.

We solve the problem using the same global list γ as with postponed reads. If a release-write is performed before a postponed one, it adds a corresponding triple to γ. Subsequently, a synchronizing acquire-read from a history record will not be performed until the corresponding triple is in the list. The only difference is that when the semantics resolves a postponed write, it does not only delete the corresponding records from γ, but also updates the synchronization front in the history record stored by the release write, therefore, "bringing it up-to-date" with the globally published stored values.

3.6 Putting it all together

As one can notice, almost every aspect of the C11 standard, outlined in Sections 3.1–3.5 requires us to enhance our semantics in one way or another. The good news are that almost all of these enhancements are *orthogonal*: they can be added to the operational model independently. For instance, one can consider a subset of C11 with RA-synchronization, relaxed and non-atomic accesses, but without accounting for release-sequences or SC-accesses.

$$
\begin{array}{rcll}
\mathtt{e} & ::= & \mathtt{x} \mid z\ (\in \mathbb{Z}) \mid \mathtt{e}_1\ \mathtt{op}\ \mathtt{e}_2 \mid \mathbf{choice}\ \mathtt{e}_1\ \mathtt{e}_2 & \\
 & & \mathbf{fst}\ \mathtt{e} \mid \mathbf{snd}\ \mathtt{e} \mid (\mathtt{e}_1,\ \mathtt{e}_2) \mid \iota & \\
\mathtt{op} & ::= & + \mid - \mid * \mid / \mid \% \mid == \mid != & \\
\iota & ::= & \ell \mid \mathtt{x} & \\
\ell & — & \text{location identifier} & \\
\mathtt{x} & — & \text{local variable} & \\
v & ::= & \ell \mid z \mid (v_1,\ v_2) & \\
 & & & \\
\mathtt{s} & ::= & \mathtt{e} \mid \mathtt{x} = \mathtt{s}_1\,;\,\mathtt{s}_2 \mid \mathbf{spw}\ \mathtt{s}_1\ \mathtt{s}_2 \mid & \\
 & & \mathbf{if}\ \mathtt{e}\ \mathbf{then}\ \mathtt{s}_1\ \mathbf{else}\ \mathtt{s}_2\ \mathbf{fi} \mid \mathbf{repeat}\ \mathtt{s}\ \mathbf{end} \mid & \\
 & & [\iota]_{\mathsf{RM}} \mid [\iota]_{\mathsf{WM}} := \mathtt{e} \mid \mathsf{cas}_{\mathsf{SM,FM}}(\iota, \mathtt{e}_1, \mathtt{e}_2) & \\
\mathsf{SRT} & ::= & \mathbf{stuck} \mid \mathbf{par}\ \mathtt{s}_1\ \mathtt{s}_2 & \\
 & & & \\
\mathsf{RM} & ::= & \mathtt{sc} \mid \mathtt{acq} \mid \mathtt{con} \mid \mathtt{rlx} \mid \mathtt{na} & \\
\mathsf{WM} & ::= & \mathtt{sc} \mid \mathtt{rel} \mid \mathtt{rlx} \mid \mathtt{na} & \\
\mathsf{SM} & ::= & \mathtt{sc} \mid \mathtt{relAcq} \mid \mathtt{rel} \mid \mathtt{acq} \mid \mathtt{con} \mid \mathtt{rlx} & \\
\mathsf{FM} & ::= & \mathtt{sc} \mid \mathtt{acq} \mid \mathtt{con} \mid \mathtt{rlx} &
\end{array}
$$

Figure 13. Syntax of statements and expressions.

$$
\frac{\xi' = \mathsf{spawn}(\mathtt{E},\ \xi)}{\langle \mathtt{E}[\mathbf{spw}\ \mathtt{s}_1\ \mathtt{s}_2], \xi\rangle \Longrightarrow \langle \mathtt{E}[\mathbf{par}\ \mathtt{s}_1\ \mathtt{s}_2], \xi'\rangle}\ \textsc{Spawn}
$$

$$
\frac{\xi' = \mathsf{join}(\mathtt{E},\ \xi)}{\langle \mathtt{E}[\mathbf{par}\ v_1\ v_2], \xi\rangle \Longrightarrow \langle \mathtt{E}[(v_1,\ v_2)], \xi'\rangle}\ \textsc{Join}
$$

Figure 14. Generic rules for spawning and joining threads

State	ξ	::=	$\langle H, \psi^{\mathtt{rd}}\rangle$
History	H	::=	$(\ell, \tau) \rightharpoonup \langle v, \sigma\rangle$
Viewfront function	$\psi^{\mathtt{rd}}$	::=	$\pi \rightharpoonup \sigma$
Viewfront	σ	::=	$\ell \rightharpoonup \tau$
Thread path	π	—	$(l \mid r)^*$
Timestamp	$\tau \in \mathbb{N}$	—	timestamp

Figure 15. States, histories and viewfronts.

4. Operational Semantics, Formally

In this section, we formally describe main components of our operational semantics for C11, starting from the definition of the language, histories and viewfronts, followed by the advanced aspects. The semantics of consume-reads is described in Appendix **??**.

4.1 Language syntax and basic reduction rules

The syntax of the core language is presented in Figure 13. The meta-variable e ranges over expressions, which might be integer numbers z, location identifiers ℓ, (immutable) local variables x, pairs, selectors and binary operations. The random **choice** operator, which non-deterministically returns one of its arguments. At the moment, arrays or pointer arithmetics are not supported.

Programs are statements, represented by terms s, most of which are standard. As customary in operational semantics, the result of a fully evaluated program is either a value v or the *run-time* **stuck** statement, which denotes the result of a program that "went wrong". For instance, it is used to indicate all kinds of undefined behavior, *e.g.*, data races on *non-atomic* operations or reading from non-initialized locations. The **spw** $\mathtt{s}_1$ $\mathtt{s}_2$, when reduced, spawns two threads with subprograms $\mathtt{s}_1$ and $\mathtt{s}_2$ respectively, emitting the *run-time* statement **par** $\mathtt{s}_1$ $\mathtt{s}_2$, which is necessary for implementing dynamic viewfront allocation for the newly forked threads, as will be described below. In our examples, we will use the parallel composition operator || for both **spw** and **par** statements.

A binding statement $\mathtt{x} = \mathtt{s}_1\,;\,\mathtt{s}_2$ implements sequential composition by means of substituting $\mathtt{s}_1$ in $\mathtt{s}_2$ for all occurrences of x. Location-manipulating statements include reading from a location ($[\iota]_{\mathsf{RM}}$), writing ($[\iota]_{\mathsf{WM}} := \mathtt{e}$), and compare-and-set on a location ι ($\mathsf{cas}_{\mathsf{SM,FM}}(\iota, \mathtt{e}_1, \mathtt{e}_2)$). These statements are annotated with order modifiers. We will sometimes abbreviate $\mathtt{r1} = [\mathtt{x}]_{\mathsf{RM}}\,;\ [\mathtt{y}]_{\mathsf{WM}} := \mathtt{r2}$ as $[\mathtt{y}]_{\mathsf{WM}} := [\mathtt{x}]_{\mathsf{RM}}$ and $\mathtt{r1} = \mathtt{s};\ \mathtt{r1}$ as $\mathtt{r1} = \mathtt{s}$.

Meta-variable ξ ranges over dynamic environments, defined further. Evaluation of a program s in the semantics starts with the initial state $\langle \mathtt{s}, \xi_{init}\rangle$, where ξ_{init} contains an empty history, and an empty viewfront for the only initial thread. The semantics is defined in reduction style [18], with most of its rules of the form

$$
\frac{\cdots}{\langle \mathtt{E}[\mathtt{s}],\ \xi\rangle \Longrightarrow \langle \mathtt{E}[\mathtt{s}'],\ \xi'\rangle}
$$

where E is a reduction context, defined as follows:

$$
\mathtt{E} \quad ::= \quad [\,] \mid \mathtt{x} = \mathtt{E}; \mathtt{s} \mid \mathbf{par}\ \mathtt{E}\ \mathtt{s} \mid \mathbf{par}\ \mathtt{s}\ \mathtt{E}
$$

If there is more than one thread currently forked, *i.e.*, the program expression contains a **par** node, its statement might be matched against E[s] in multiple possible ways non-deterministically.

The core rules of our semantics, involving non-memory operations, are standard and are presented in Appendix B. The only interesting rules are for spawning and joining threads (Figure 14), as they alter the thread-related information in the environment (*e.g.*, the viewfronts). The exact shape of these rules depends on the involved concurrency aspects, which define the meta-functions spawn and join.

4.2 Histories and Viewfronts

In its simplest representation, the program environment ξ is a pair, whose components are a history H and per-thread viewfront function $\psi^{\mathtt{rd}}$, defined in Figure 15.

A history H is a partial function from location identifiers ℓ and timestamps τ to pairs of a stored value v and a synchronization front σ. Further aspects of our semantics feature different kinds of fronts, but all of them have the same shape, mapping memory locations to timestamps. Per-thread viewfront function $\psi^{\mathtt{rd}}$ maps *thread paths* (π) to viewfronts. A thread path is a list of directions $(l \mid r)$, which shows how to get to the thread subexpression inside a program statement tree through the **par** nodes: it uniquely identifies a thread in a program statement. We use an auxiliary function path in the rules to calculate a path from an evaluation context E.

Once threads are spawn, they inherit a viewfront of their parent thread, hence the simplest spawn function is defined as follows:

$$
\mathsf{spawn}(\mathtt{E},\ \langle \mathtt{s}, \psi^{\mathtt{rd}}\rangle) \triangleq \langle \mathtt{s}, \psi^{\mathtt{rd}}[\pi\ l \mapsto \sigma_{\mathtt{rd}},\ \pi\ r \mapsto \sigma_{\mathtt{rd}}]\rangle
$$

where $\pi = \mathsf{path}(\mathtt{E})$, and $\sigma_{\mathtt{rd}} = \psi^{\mathtt{rd}}(\pi)$. When threads join, their parent thread gets a viewfront, which is the least upper bound (join) of subthread viewfronts:

$$
\mathsf{join}(\mathtt{E},\ \langle \mathtt{s}, \psi^{\mathtt{rd}}\rangle) = \langle \mathtt{s}, \psi^{\mathtt{rd}}[\pi \mapsto \sigma^l_{\mathtt{rd}} \sqcup \sigma^r_{\mathtt{rd}}]\rangle
$$

where $\pi = \mathsf{path}(\mathtt{E})$, $\sigma^l_{\mathtt{rd}} = \psi^{\mathtt{rd}}(\pi\ l)$, and $\sigma^r_{\mathtt{rd}} = \psi^{\mathtt{rd}}(\pi\ r)$, thus, synchronizing the children threads' views.

We can now define the first class of "wrong" behaviors, corresponding to reading from non-initialized locations (Figure 16). The rules are applicable in the case when a thread tries to read from a location, which it knows nothing about, *i.e.*, its viewfront is not yet defined for the location, making it uninitialized from the thread's point of view. This condition is also satisfied if the location is not initialized at all, *i.e.*, it has no corresponding records in the history.

4.3 Release/Acquire synchronization

The reduction rules for release-write and acquire-read are given in Figure 17. A release-write augments the history with a new entry $(\ell, \tau) \mapsto (v, \sigma)$, where v is a value argument of the write, and σ is a synchronization front, which now might be retrieved by the threads reading from the new history entry. The stored front σ is the same

$$\frac{\xi = \langle H, \psi^{\mathtt{rd}}\rangle \quad \pi = \mathsf{path}(\mathtt{E}) \quad \sigma_{\mathtt{rd}} = \psi^{\mathtt{rd}}(\pi) \quad \sigma_{\mathtt{rd}}(\ell) = \bot}{\langle \mathtt{E}[[\ell]_{\mathtt{RM}}], \xi\rangle \Longrightarrow \langle \mathbf{stuck}, \xi_{init}\rangle}\ \textsc{Read-Uninit}$$

$$\frac{\xi = \langle H, \psi^{\mathtt{rd}}\rangle \quad \pi = \mathsf{path}(\mathtt{E}) \quad \sigma_{\mathtt{rd}} = \psi^{\mathtt{rd}}(\pi) \quad \sigma_{\mathtt{rd}}(\ell) = \bot}{\langle \mathtt{E}[\mathsf{cas}_{\mathsf{SM},\mathsf{FM}}(\ell, \mathtt{e}_1, \mathtt{e}_2)], \xi\rangle \Longrightarrow \langle \mathbf{stuck}, \xi_{init}\rangle}\ \textsc{CAS-Uninit}$$

Figure 16. Rules for reading from an uninitialized location.

$$\frac{\begin{array}{c}\xi = \langle H, \psi^{\mathtt{rd}}\rangle \quad \pi = \mathsf{path}(\mathtt{E}) \quad \tau = \mathit{Next}\tau(H, \ell)\\ \sigma_{\mathtt{rd}} = \psi^{\mathtt{rd}}(\pi) \quad \sigma = \sigma_{\mathtt{rd}}[\ell \mapsto \tau]\\ \xi' = \langle H[(\ell, \tau) \mapsto (v, \sigma)], \psi^{\mathtt{rd}}[\pi \mapsto \sigma]\rangle\end{array}}{\langle \mathtt{E}[[\ell]_{\mathtt{rel}} := v], \xi\rangle \Longrightarrow \langle \mathtt{E}[v], \xi'\rangle}\ \textsc{WriteRel}$$

$$\frac{\begin{array}{c}\xi = \langle H, \psi^{\mathtt{rd}}\rangle \quad \pi = \mathsf{path}(\mathtt{E}) \quad H(\ell, \tau) = (v, \sigma)\\ \sigma_{\mathtt{rd}} = \psi^{\mathtt{rd}}(\pi) \quad \sigma_{\mathtt{rd}}(\ell) \le \tau\\ \xi' = \langle H, \psi^{\mathtt{rd}}[\pi \mapsto \sigma_{\mathtt{rd}} \sqcup \sigma]\rangle\end{array}}{\langle \mathtt{E}[[\ell]_{\mathtt{acq}}], \xi\rangle \Longrightarrow \langle \mathtt{E}[v], \xi'\rangle}\ \textsc{ReadAcq}$$

Figure 17. Reduction rules for release/acquire atomics.

$$\frac{\xi = \langle \ldots, \sigma^{\mathtt{sc}}\rangle \quad \ldots \quad \xi' = \langle \ldots, \sigma^{\mathtt{sc}}[\ell \mapsto \tau]\rangle}{\langle \mathtt{E}[[\ell]_{\mathtt{sc}} := v], \xi\rangle \Longrightarrow \langle \mathtt{E}[v], \xi'\rangle}\ \textsc{WriteSC}$$

$$\frac{\xi = \langle \ldots, \sigma^{\mathtt{sc}}\rangle \quad \ldots \quad \mathit{max}(\sigma_{\mathtt{rd}}(\iota), \sigma^{\mathtt{sc}}(\iota)) \le \tau}{\langle \mathtt{E}[[\ell]_{\mathtt{sc}}], \xi\rangle \Longrightarrow \langle \mathtt{E}[v], \xi'\rangle}\ \textsc{ReadSC}$$

Figure 18. Reduction rules for SC atomics.

as the viewfront of the writer thread after having stored the value, *i.e.*, featuring updated ℓ-entry with the new timestamp τ.

An acquire-read is more interesting. It *non-deterministically* chooses, from the global history, an entry $(\ell, \tau) \mapsto (v, \sigma)$, with a timestamp τ which is at least a new as the timestamp τ' for the corresponding location ℓ in the thread's viewfront $\sigma_{\mathtt{rd}}$ (*i.e.*, $\tau' \le \tau$). The value v replaces the read expression inside the context $\mathtt{E}$, and the thread's local viewfront $\sigma_{\mathtt{rd}}$ is updated via the retrieved synchronization front σ. The RA-CAS operations (rules omitted for brevity) behave similarly with only difference: a successful CAS reads from the *latest* entry in the history.

4.4 SC operations

To account for SC operations, we augment ξ with sc-front:

$$\xi \quad ::= \quad \langle \ldots, \sigma^{\mathtt{sc}}\rangle$$

that maps each location to a timestamp of the *latest* entry in the location history, which has been added by a SC-write.

The SC operations update the history and local/stored fronts similarly to RA atomics. In addition, an SC-write updates sc-front, and an SC-read introduces an additional check for the timestamp τ, taking max of two viewfronts, as defined in Figure 18.[2] That is, the rule ReadSC ensures that an SC-read gets a history entry, which is not older than the one added by the last SC write to the location.

4.5 Non-atomic operations

For non-atomic accesses we augment ξ with na-front:

$$\xi \quad ::= \quad \langle \ldots, \sigma^{\mathtt{na}}\rangle$$

Similarly to sc-front, it maps a location to the latest corresponding na-entry in the history, and it is updated by NA-writes.

[2] Read and write rules only depict difference with the release/acquire ones.

$$\frac{\begin{array}{c}\xi = \langle H, \psi^{\mathtt{rd}}, \ldots, \sigma^{\mathtt{na}}\rangle \quad \ldots\\ \sigma_{\mathtt{rd}}(\ell) \equiv \mathsf{LastTS}(H, \ell) \quad \sigma = \sigma_{\mathtt{rd}}[\ell \mapsto \tau]\\ \xi' = \langle H[(\ell, \tau) \mapsto (v, ())], \psi^{\mathtt{rd}}[\pi \mapsto \sigma], \ldots, \sigma^{\mathtt{na}}[\ell \mapsto \tau]\rangle\end{array}}{\langle \mathtt{E}[[\ell]_{\mathtt{na}} := v], \xi\rangle \Longrightarrow \langle \mathtt{E}[v], \xi'\rangle}\ \textsc{WriteNA}$$

$$\frac{\begin{array}{c}\xi = \langle H, \psi^{\mathtt{rd}}, \ldots, \sigma^{\mathtt{na}}\rangle \quad \ldots\\ \tau = \mathsf{LastTS}(H, \ell) \quad \tau \equiv \sigma_{\mathtt{rd}}(\ell) \quad H(\ell, \tau) = (v, \sigma)\end{array}}{\langle \mathtt{E}[[\ell]_{\mathtt{na}}], \xi\rangle \Longrightarrow \langle \mathtt{E}[v], \xi\rangle}\ \textsc{ReadNA}$$

$$\frac{\xi = \langle H, \psi^{\mathtt{rd}}, \ldots, \sigma^{\mathtt{na}}\rangle \quad \ldots \quad \sigma_{\mathtt{rd}}(\ell) \neq \mathsf{LastTS}(H, \ell)}{\langle \mathtt{E}[[\ell]_{\mathtt{na}}], \xi\rangle \Longrightarrow \langle \mathbf{stuck}, \xi\rangle}\ \textsc{ReadNA-stuck1}$$

$$\frac{\xi = \langle H, \psi^{\mathtt{rd}}, \ldots, \sigma^{\mathtt{na}}\rangle \quad \ldots \quad \sigma_{\mathtt{rd}}(\ell) < \sigma^{\mathtt{na}}(\ell)}{\langle \mathtt{E}[[\ell]_{\mathtt{RM}}], \xi\rangle \Longrightarrow \langle \mathbf{stuck}, \xi\rangle}\ \textsc{ReadNA-stuck2}$$

Figure 19. Reduction rules for non-atomics.

$$\frac{\xi = \langle H, \psi^{\mathtt{rd}}, \ldots\rangle \quad \ldots \quad \xi' = \langle H, \psi^{\mathtt{rd}}[\pi \mapsto \sigma_{\mathtt{rd}}[\ell \mapsto \tau]], \ldots\rangle}{\langle \mathtt{E}[[\ell]_{\mathtt{rlx}}], \xi\rangle \Longrightarrow \langle \mathtt{E}[v], \xi'\rangle}\ \textsc{ReadRlx}$$

$$\frac{\begin{array}{c}\xi = \langle H, \psi^{\mathtt{rd}}, \ldots, \psi^{\mathtt{wr}}\rangle \quad \ldots\\ \tau_{rel} = \psi^{\mathtt{wr}}(\pi)(\ell) \quad (_, \sigma_{\mathsf{sync}}) = H(\ell, \tau_{rel})\\ \xi' = \langle H[(\ell, \tau) \mapsto (v, \sigma_{\mathsf{sync}}[\ell \mapsto \tau])],\\ \psi^{\mathtt{rd}}[\pi \mapsto \sigma_{\mathtt{rd}}[\ell \mapsto \tau]], \ldots, \psi^{\mathtt{wr}}\rangle\end{array}}{\langle \mathtt{E}[[\ell]_{\mathtt{rlx}} := v], \xi\rangle \Longrightarrow \langle \mathtt{E}[v], \xi'\rangle}\ \textsc{WriteRlx}$$

$$\frac{\begin{array}{c}\xi = \langle H, \psi^{\mathtt{rd}}, \ldots, \psi^{\mathtt{wr}}\rangle \quad \ldots\\ \sigma_{\mathtt{wr}} = \psi^{\mathtt{wr}}(\pi) \quad \sigma'_{\mathtt{wr}} = \sigma_{\mathtt{wr}}[\ell \mapsto \tau]\\ \xi' = \langle \ldots, \psi^{\mathtt{wr}}[\pi \mapsto \sigma'_{\mathtt{wr}}]\rangle\end{array}}{\langle \mathtt{E}[[\ell]_{\mathtt{rel}} := v], \xi\rangle \Longrightarrow \langle \mathtt{E}[v], \xi'\rangle}\ \textsc{WriteRel'}$$

Figure 20. Reduction rules for relaxed atomics.

The real purpose of na-front is to detect *data races* involving NA-operations as defined in Figure 19. When a thread performs NA-write or NA-read (see WriteNA, ReadNA rules), it must be aware of the latest stored record of the location (*i.e.*, it should match the timestamp in its local front $\sigma_{\mathtt{rd}}$). Violating this side condition is condemned to be a data race and leads to undefined behavior (ReadNA-stuck1). In addition, if a thread performs *any* write to or read from a location, it should be aware of the latest NA-record to the location (see the side condition $\sigma_{\mathtt{rd}}(\ell) < \sigma^{\mathtt{na}}(\ell)$ in the rule ReadNA-stuck2).[3] The stuck-cases reflect the cases when a write or a read is in data race with the last NA-write to the location.

Unlike release or *SC*-writes, NA-writes do not store a front to the history entry, as they cannot be used for synchronization. A similar fact holds for NA-reads: they do not get a stored front from the history entry, upon reading from it.

4.6 Release-sequences and write-fronts

Relaxed reads do not update their thread's viewfront with a synchronization front from the history (see ReadRlx rule in Figure 20). At this stage, their support does not require augmenting the state.

An additional instrumentation is required, though, to encode release sequences. As discussed in Section 3.5, an acquire-read, when reading the result of a relaxed write, might get synchronized with a release-write to the same location performed earlier by the same writer thread. To account for this, we introduce per-thread

[3] Rules WriteNA-stuck1 and WriteNA-stuck2 are similar and can be found in Appendix B.

write-front function $\psi^{\mathtt{wr}}$ as an environment component:

$$\xi \quad ::= \quad \langle \ldots, \psi^{\mathtt{wr}} \rangle$$

It is similar to the viewfront function $\psi^{\mathtt{rd}}$, but it stores a timestamp of the last release-write to a location by the thread. Specifically, when a thread performs a relaxed write W (see rule WriteRlx in Figure 20), it checks if there was a release-write W' performed by it earlier, takes a synchronization front σ_{sync} from the history entry, added by W', and stores it as the synchronization front in the *new* history entry.[4] Additionally, we need to modify the old rules WriteRel, WriteSC, CAS-Rel, *etc*., so they update $\psi^{\mathtt{wr}}$ correspondingly (*e.g.*, see rule WriteRel' in Figure 20).

We also need to change our meta-functions, in order to account for the $\psi^{\mathtt{wr}}$ component of the state environment:

$$\begin{array}{rcl} \mathsf{spawn}(\mathsf{E},\ \langle \ldots, \psi^{\mathtt{wr}} \rangle) & = & \langle \ldots, \psi^{\mathtt{wr}}[\pi\ l \mapsto \bot,\ \pi\ r \mapsto \bot] \rangle \\ \mathsf{join}(\mathsf{E},\ \langle \ldots, \psi^{\mathtt{wr}} \rangle) & = & \langle \ldots, \psi^{\mathtt{wr}}[\pi \mapsto \bot] \rangle \end{array}$$

Subthreads do not inherit write-fronts upon spawning, and a parent thread does not inherit the joined one, since the described synchronization effects via relaxed writes and release sequences do not propagate through spawn/join points according to the model [7].

4.7 Postponed operations and speculations

To support postponed operations (or, equivalently, speculative executions)we instrument the state with two additional components:

$$\xi \quad ::= \quad \langle \ldots, \varphi, \gamma \rangle$$

The main one, φ, is a function that maps a thread path π to a per-thread hierarchical buffer α of postponed operations β:

$$\begin{array}{lll} \varphi & ::= & \pi \rightharpoonup \alpha \\ \alpha & ::= & \beta^* \\ \beta & ::= & \mathit{read}\langle \mathsf{x}, \iota, \mathsf{RM} \rangle \mid \mathit{write}\langle \mathsf{x}, \iota, \mathsf{WM}, \mathsf{e} \rangle \mid \mathit{bind}\langle \mathsf{x}, \mathsf{e} \rangle \mid \mathit{if}\langle \mathsf{x}, \mathsf{e}, \alpha, \alpha \rangle \end{array}$$

Each operation β is uniquely identified by its symbolic value x. Read entries contain a (possibly unresolved) location ι to read from as well as a read modifier RM. Write entries additionally contain an expression e to be stored to the location. *Bind* entries are used to postpone calculation of an expression depending on a symbolic value, making it possible to postpone the reads as follows:

$$\mathtt{r1} = [\mathtt{x}]_{rlx};\ \mathtt{r2} = \mathtt{r1} + 1;\ \ldots$$

Both reads r1 and bind r2 might be postponed, so the second statement will not "trigger" evaluation of the first one. *If*-entries have a conditional expression e and two *subbuffers* α representing operations speculatively put to the buffer under *then* and *else* branches. To represent speculation under (possibly nested) **if** statement we define an if-specialized reduction context $\mathsf{E}\alpha$ as follows:

$$\begin{array}{lll} \mathsf{E}\alpha & ::= & [\ ] \mid \mathsf{x} = \mathsf{E}\alpha; \mathsf{s} \mid \mathtt{if}\ \mathsf{x}\ \mathtt{then}\ \mathsf{E}\alpha\ \mathtt{else}\ \mathsf{s}\ \mathtt{fi} \mid \\ & & \mathtt{if}\ \mathsf{x}\ \mathtt{then}\ \mathsf{s}\ \mathtt{else}\ \mathsf{E}\alpha\ \mathtt{fi} \end{array}$$

where the symbolic value x in the condition is the same as in the corresponding buffer entry $\mathit{if}\langle \mathsf{x}, \alpha_1, \alpha_2 \rangle$. The list of symbolic values from conditions of the context can be used to uniquely identify an operation buffer inside an hierarchical per-thread buffer α.

The list γ encodes *Acquire-Read Restrictions* by containing triples $\langle \ell, \tau, \mathsf{x} \rangle$, forbidding to acquire-read from (ℓ, τ) until x is not resolved.

During a thread execution, any read, write, or bind operation (including those under not fully reduced if-branches) can be postponed by the semantics by adding a corresponding record into the matching subbuffer of the thread buffer α. For the sake of brevity we discuss only write rules here; other rules can be found in the appendix and in our implementation.

The postpone-rules append operation records into the corresponding buffer φ. An operation to be postponed can be nested under a not fully reduced if statement.

$$\frac{\begin{array}{c} \xi = \langle \ldots, \varphi, \gamma \rangle \quad \ldots \\ \mathsf{x} \text{ is fresh symbolic variable} \quad \alpha = \varphi(\pi) \\ \xi' = \langle \ldots, \varphi[\pi \mapsto \mathsf{append}(\alpha, \mathsf{E}\alpha, \mathit{write}\langle \mathsf{x}, \iota, \mathsf{WM}, \mathsf{e} \rangle)], \gamma \rangle \end{array}}{\langle \mathsf{E}[\mathsf{E}\alpha[[\iota]_{\mathsf{WM}} := \mathsf{e}]], \xi \rangle \Longrightarrow \langle \mathsf{E}[\mathsf{E}\alpha[\mathsf{x}]], \xi' \rangle}\ \textsc{Write-Postpone}$$

After postponing, a write-record has to be resolved eventually:

$$\frac{\begin{array}{c} \xi = \langle H, \ldots, \psi^{\mathtt{wr}}, \varphi, \gamma \rangle \quad \ldots \\ \mathit{write}\langle \mathsf{x}, \ell, \mathsf{WM}, v \rangle \in \varphi(\pi) \text{ and not in conflict} \\ \varphi' = \mathsf{remove}(\varphi, \pi, \mathsf{x}) \\ \gamma' = \mathsf{updateDep}(\mathsf{x}, \mathsf{WM}, \psi^{\mathtt{wr}}, \ell, \tau, \gamma, \varphi(\pi)) \\ H' = \mathsf{updateSync}(\mathsf{x}, \mathsf{WM}, \sigma, \gamma, H) \\ \xi' = \langle H'[(\ell, \tau) \mapsto (v, \sigma)], \ldots, \varphi'[v/\mathsf{x}], \gamma' \rangle \end{array}}{\langle \mathsf{s}, \xi \rangle \Longrightarrow \langle \mathsf{s}[v/\mathsf{x}], \xi' \rangle}\ \textsc{Write-Resolve}$$

We point out several important side conditions of the rule. The write must be in the top-level buffer $\varphi(\pi)$, and there must be no operation *before* it in the buffer, which is in conflict with it, *e.g.*, an acquire-read or a write to the same location (line 2). In line 4, updateDep updates γ as follows. First, it removes from γ all entries that mention the symbolic value x, $\langle \ell', \tau', \mathsf{x} \rangle$, since the write is resolved. Second, for specific postponed operations x', it adds $\langle \ell, \tau, \mathsf{x}' \rangle$ entries to γ, blocking acquire-reads from the newly created history entry (ℓ, τ) until x' are resolved. These are the operations that can affect an acquire-read from (ℓ, τ): *(i)* ones related by γ to $(\ell, \psi^{\mathtt{wr}}(\pi, \ell))$ (*i.e.*, a record of the last release-write), and *(ii)*, if the resolved write x is a release one, operations, that precede the write in the thread buffer $\varphi(\pi)$ or are unresolved writes observed by thread (*i.e.*, elements of $\omega(\pi)$). In line 6, updateSync updates synchronization fronts in history records (ℓ', τ') related to x by γ.

Duplicated non-conflicting writes from nested buffers can be (non-deterministically) promoted to an upper-level:

$$\frac{\begin{array}{c} \xi = \langle H, \psi^{\mathtt{rd}}, \ldots, \varphi, \gamma \rangle \quad \ldots \\ \mathit{if}\langle \mathsf{x}'', \mathsf{e}, \alpha_1, \alpha_2 \rangle \text{ is inside } \varphi(\pi) \\ \mathit{write}\langle \mathsf{x}, \ell, \mathsf{WM}, v \rangle \in \alpha_1 \quad \mathit{write}\langle \mathsf{x}', \ell, \mathsf{WM}, v \rangle \in \alpha_2 \\ \mathsf{x}, \mathsf{x}' \text{ are not in conflict in } \alpha_1, \alpha_2 \\ \varphi' = \mathsf{promote}(\mathsf{x}, \mathsf{x}', \varphi) \quad \gamma' = \gamma[\mathsf{x}'/\mathsf{x}] \end{array}}{\langle \mathsf{s}, \xi \rangle \Longrightarrow \langle \mathsf{s}[\mathsf{x}'/\mathsf{x}], \xi' \rangle}\ \textsc{Write-Promote}$$

Line 3 of the rule's premise requires two identical writes to be present in the "sibling" buffers. In line 4, promote removes the writes from α_1 and α_2, and puts $\mathit{write}\langle \mathsf{x}, \ell, \mathsf{WM}, v \rangle$ in the parent buffer before $\mathit{if}\langle \mathsf{x}'', \mathsf{e}, \alpha_1', \alpha_2' \rangle$.

The rule for initialization of speculative execution of branches of an **if**-statement adds $\mathit{if}\langle \mathsf{x}, \mathsf{e}, \langle\rangle, \langle\rangle \rangle$ into α, similarly to postponing a write, and replaces the condition with a symbolic value:

$$\frac{\begin{array}{c} \xi = \langle \ldots, \varphi, \gamma \rangle \quad \ldots \\ \mathsf{e} \text{ depends on an unresolved symbolic value} \\ \mathsf{x} \text{ — fresh symbolic variable} \quad \alpha = \varphi(\pi) \\ \xi' = \langle \ldots, \varphi[\pi \mapsto \mathsf{append}(\alpha, \mathsf{E}\alpha, \mathit{if}\langle \mathsf{x}, \mathsf{e}, \langle\rangle, \langle\rangle \rangle)], \gamma \rangle \end{array}}{\begin{array}{c} \langle \mathsf{E}[\mathsf{E}\alpha[\mathtt{if}\ \mathsf{e}\ \mathtt{then}\ \mathsf{s}_1\ \mathtt{else}\ \mathsf{s}_2\ \mathtt{fi}]], \xi \rangle \Longrightarrow \\ \langle \mathsf{E}[\mathsf{E}\alpha[\mathtt{if}\ \mathsf{x}\ \mathtt{then}\ \mathsf{s}_1\ \mathtt{else}\ \mathsf{s}_2\ \mathtt{fi}]], \xi' \rangle \end{array}}\ \textsc{If-Speculation-Init}$$

Finally, upon resolving all symbolic values in the condition e of an **if**-statement, the statement itself can be reduced:

$$\frac{\begin{array}{c} \xi = \langle \ldots, \varphi, \gamma \rangle \quad \ldots \quad \mathit{if}\langle \mathsf{x}, z, \alpha_1, \alpha_2 \rangle \in \varphi(\pi) \quad z \neq 0 \\ \varphi' = \varphi[\mathit{if}\langle \mathsf{x}, z, \alpha_1, \alpha_2 \rangle / \alpha_1] \quad \xi' = \langle \ldots, \varphi', \gamma \rangle \end{array}}{\langle \mathsf{E}[\mathsf{E}\alpha[\mathtt{if}\ \mathsf{x}\ \mathtt{then}\ \mathsf{s}_1\ \mathtt{else}\ \mathsf{s}_2\ \mathtt{fi}]], \xi \rangle \Longrightarrow \langle \mathsf{E}[\mathsf{E}\alpha[\mathsf{s}_1]], \xi' \rangle}\ \textsc{If-Resolve-True}$$

5. Implementation and Evaluation

We implemented our semantics in PLT Redex,[5] a framework on top of the Racket programming language [17, 25]. The implementation of the core language definitions (Sections 4.1–4.2) is 2070 LOC,

[4] If $H(\ell, \tau_{rel}) = \bot$ then $\sigma_{\mathsf{sync}} = \bot$.

[5] The sources are available as supplementary material for the paper.

with various C11 concurrency aspects (Sections 4.3–4.7) implemented on top of them, in 1310 LOC. Implementation of litmus tests (Section 5) and case studies (Section 6) took 3130 LOC.

Evaluation via Litmus Tests To ensure the adequacy of our semantics with respect to the C++11 standard [2] and gain confidence in its implementation, we evaluated it on a number of litmus test programs from the literature. For each test, we encoded the set of expected results and checked, via extensive state-space enumeration, provided by PLT Redex, that these are the only outcomes produced.

Figure 21 provides a table, relating specific litmus tests from the literature [7, 9, 26, 29, 45] to the relevant aspects of our semantics from Section 4, required in order to support their desired behavior. All tests mentioned before in the paper are presented in the table. Since there is no common naming conventions for litmus tests in a high-level language, making consistent appearance in related papers, we supplied ours with meaningful names, grouping them according to the behavioral pattern they exercise (*e.g.*, message-passing, store buffering, *etc.*). Exact definitions of the test program and descriptions of their behaviors can be found in Appendix A.

All tests within the same group have a similar structure but differ in memory access modifiers. The columns Hst–JN in Figure 21 show, which semantic aspects a test requires for its complete and correct execution. The last column indicates, whether the test's behavior in our semantics matches fully its outcome according to the C11 standard or not. Below, we discuss the tests that behave differently in the C11 standard and in our semantics.

Discrepancies with the C11 standard The combining relaxed and acquire-writes, LB_{acq+rlx,acq+rlx+join}, in order to adhere to the "canonical" C11 behavior, require an ability to do an acquire-read and a subsequent relaxed write out-of-order. Even though the relaxed behavior of this kind is not supported by our semantics, it is not observable under sound compilation schemes of acquire-read to the major architectures [1]: *(i)* load buffering is not observable on x86 in general [42], *(ii)* all barriers (*sync*, *lwsync*, *ctrl+isync*) forbid reorderings of read and write on Power [4], *(iii)* as well as barriers (*dmb sy*, *dmb ld*, *ctrl+isb*) on ARM [19].

Our semantics rules out OTA behaviors (tests OTA_{lb, if}), which are considered to be an issue of the standard [5, 8, 38].

Test name	VF	WF	SCF	NAF	PO	ARR	CR	JN	**C11**
Store Buffering (SB), §A.1									
rel+acq	✓								✓
sc	✓		✓						✓
sc+rel	✓		✓						✓
sc+acq	✓		✓						✓
Load Buffering (LB), §A.2									
rlx	✓				✓				✓
rel+rlx	✓				✓				✓
acq+rlx	✓				✓				✗
rel+acq+rlx	✓				✓	✓			✓
rlx+use	✓				✓				✓
rlx+let	✓				✓				✓
rlx+join	✓				✓			✓	✓
rel+rlx+join	✓				✓			✓	✓
acq+rlx+join	✓				✓			✓	✗
Message passing (MP), §A.3									
rlx+na	✓			✓					✓
rel+rlx+na	✓			✓					✓
rlx+acq+na	✓			✓					✓
rel+acq+na	✓			✓		✓			✓
rel+acq+na+rlx(_2)	✓	✓		✓		✓			✓
con+na(_2)	✓			✓			✓		✓
cas+rel+acq+na	✓			✓		✓			✓
cas+rel+rlx+na	✓			✓					✓
Coherence of Read-Read (CoRR), §A.4									
rlx	✓								✓
rel+acq	✓								✓
Independent Reads of Independent Writes (IRIW), §A.5									
rlx	✓								✓
rel+acq	✓								✓
sc	✓		✓						✓
Write-to-Read Causality (WRC), §A.6									
rlx	✓								✓
rel+acq	✓								✓
cas+rel	✓					✓			✓
cas+rlx	✓								✓
Out-of-Thin-Air (OTA), §A.7									
lb	✓				✓				✗
if	✓				✓				✗
Write Reorder (WR), §A.8									
rlx	✓				✓				✓
rlx+rel	✓				✓	✓			✓
rel	✓				✓	✓			✓
Speculative Execution (SE), §A.9									
simple	✓				✓				✓
prop	✓				✓				✓
nested	✓				✓				✓
Locks, §A.10									
Dekker	✓			✓					✓
Cohen [45]	✓			✓					✓

Figure 21. Litmus tests (Appendix A) and corresponding semantic aspects of our framework: viewfronts (VF, §2.1), write-fronts (WF, §4.6), SC-fronts (SCF, §3.1), non-atomic fronts (NAF, §3.2), postponed operations (PO, §4.7), acquire read restrictions (γ) (ARR, §4.7), consume-reads (CR), joining threads with non-empty operation buffers (JN, §3.4). The column **C11** indicates whether the behavior is coherent with the C11 standard.

6. Case Study: Read-Copy-Update

We showcase our implementation by testing and debugging a Read-Copy-Update structure (RCU) [31, 34] and its client programs.

6.1 RCU: background and implementation

Read-Copy-Update is a standard way to implement non-blocking sharing of a linked data structure (*e.g.*, list or a tree) between single writer and multiple readers, running concurrently. For our purposes, we focus on RCU for a singly linked list, implemented via *Quiescent State Based Reclamation* (QSBR) technique [14]. The central idea of RCU is the way the writer treats nodes of the linked structure. Specifically, instead of in-place modification of a list node, the writer creates a *copy* of it, modifies the copy, *updates* the link to the node, making the older version inaccessible, and then waits until *all* readers stop using the older version, so it could be reclaimed. The crux of the algorithm's correctness is a fine-grained synchronization between the writer and the readers: the writer updates the link via release-write, and the readers must traverse the list using an acquire-read for dereferencing its nodes, ensuring that readers will not observe partially modified nodes.

A QSBR RCU implementation and its client program are shown in Figure 22. The first line in the top of the figure initializes thread counters, which are used by the reader threads to signal if they use the list or not (*i.e.*, they are in a quiescent state), and a pointer to the list (`lhead`), which is going to be shared between the threads. Next, three threads are spawned: a writer and two readers. The writer thread on the left appends 1, 10, and 100 to the list.[6] A call to `append` creates a new node and adds a link from the current last node to the new one via relaxed write to `ltail`. The updating write to the last node `[rt]` is a release one, guaranteeing that a reader thread, which might be observing the added node pointer via an acquire-read in concurrent `traverse` call, will become aware of the value and

[6] In the absence of implemented allocation, the example uses the fixed locations `a`, `b`, and `c` for storing nodes of the list.

Program:

```
[cw]_na := 0; [cr1]_na := 0; [cr2]_na := 0; [lhead]_na := null;

[a]_rlx := (1, null);
[ltail]_na := a;
[lhead]_rel := a;
append(b, 10 , ltail);
append(c, 100, ltail);
updateSecondNode(d, 1000)

||

[sum11]_na := 0;
rcuOnline (cw, cr1);
traverse  (lhead, cur1, sum11)
rcuOffline(cw, cr1);

[sum12]_na := 0;
rcuOnline (cw, cr1);
traverse  (lhead, cur1, sum12)
rcuOffline(cw, cr1);

r11 = [sum11]_na;
r12 = [sum12]_na

||

[sum21]_na := 0;
rcuOnline (cw, cr2);
traverse  (lhead, cur2, sum21)
rcuOffline(cw, cr2);

[sum22]_na := 0;
rcuOnline (cw, cr2);
traverse  (lhead, cur2, sum22);
rcuOffline(cw, cr2);

r21 = [sum21]_na;
r22 = [sum22]_na
```

Functions:

```
append(loc, value, ltail) ≜
  [loc]_rlx := (value, null);
  rt  = [ltail]_na;
  rtc = [rt]_rlx;
  [rt]_rel    := (fst rtc, loc);
  [ltail]_na := loc
updateSecondNode(loc, value) ≜
  r1  = [lhead]_rlx;
  r1c = [r1]_rlx;
  r2  = snd r1c;
  r2c = [r2]_rlx;
  r3  = snd r2c;
  [loc]_rel := (value, r3);
  [r1]_rel  := (fst r1c, loc);
  sync(cw, cr1, cr2);
  delete r2
```

```
traverse(lhead, curNodeLoc, resLoc) ≜
  rh = [lhead]_acq;
  [curNodeLoc]_na := rh;
  repeat
    rCurNode = [curNodeLoc]_na;
    if (rCurNode != null)
    then rNode = [rCurNode]_acq;
         rRes = [resLoc]_na;
         rVal = fst rNode;
         [resLoc]_na := rVal + rRes;
         [curNodeLoc]_na := snd rNode;
         0
    else 1
    fi
  end
```

```
sync(cw, cr1, cr2) ≜
  rcw  = [cw]_rlx;
  rcwn = rcw + 2;
  [cw]_rel := rcwn;
  syncWithReader(rcwn, cr1);
  syncWithReader(rcwn, cr2)

syncWithReader(rcwn, cr) ≜
  repeat [cr]_acq >= rcwn end

rcuOnline(cw, cr) ≜
  [cr]_rlx := [cw]_acq + 1

rcuOffline(cw, cr) ≜
  [cr]_rel := [cw]_rlx
```

Figure 22. The QSBR RCU implementation. Fragments in gray boxes are later removed for testing purposes (Section 6.3).

link data, stored to the node. This release/acquire synchronization eliminates a potential data race.

At the end, the writer thread changes the second value in the list, 10, to 1000. In the corresponding `updateSecondNode` routine, first five commands get pointers to first, second, and third list nodes. The next two commands create a new node with `value` stored in it, and update the corresponding link in the previous node (*i.e.*, the first one). By executing `sync(cw, cr1, cr2)`, the writer checks that the reader threads no longer use an older version of the list (with 10), so, once the check succeeds, the old node can be reclaimed.

The reader threads calculate two times the sum of the list's elements by traversing the list. Before and after each traversal they call `rcuOnline` and `rcuOffline` routines respectively to signal to the writer about their state.

6.2 Additional infrastructure for testing RCU

Having an executable operational semantics gives us a possibility to run a dynamic analysis of the RCU and its client, exercising all possible executions. Realistically, running such an analysis would take forever, because of the size of the program state-space. The state-space explosion is because of the following three reasons:

1. Non-determinism due to concurrent thread scheduling;
2. Resolution of postponed operations;
3. Loading any value, which is *newer* than the one in the thread's viewfront representation of the reading location.

Indeed, our semantics accounts for all combinations of the factors above, exploring all possible execution traces.

Randomized semantics In order to make dynamic analysis practically feasible, we implemented a semantics, which non-deterministically chooses a random path in the program state space of the original semantics. It does so by applying semantic rules to the current state getting a set of new states, checks if there is a **`stuck`** state, and randomly chooses the next state from the set. The presence of the randomized semantics makes possible to implement property-based of testing of executions [22].

Deallocation As an additional aspect, we added **`delete`** operator to the language for reclaiming retired nodes in RCU, and extended the state by a global list of reclaimed locations. That is, if a location is added to the "retired" list, any read or write on it will lead to the **`stuck`** state, indicating accessing a deallocated pointer.

6.3 Testing and debugging the RCU implementation

We can now run some random tests on our RCU implementation and see whether it meets the basic safety requirements. In particular, we can check that no matter what path is being exercised by the randomized semantics, the execution of the program does not get stuck. In our experience, this correctness condition held for the implementation in Figure 22 for all test runs.

Next, we intentionally introduced a synchronization bug into our implementation. In particular, we removed from the implementation `syncWithReader` loops (see grayed code fragments), which were used to synchronize the writer with the readers. Additionally, we considered the following correctness criteria: *(i)* the values of `r11`, `r12`, `r21`, and `r22` must be in $\{0, 1, 11, 111, 1101\}$: this guarantees that the list is read correctly, and, *(ii)* it should be the case that by the end $\texttt{r11} \le \texttt{r12} \wedge \texttt{r21} \le \texttt{r22}$, *i.e.*, second list traversals see the list at least as up to date as the first ones.

We ran the test twenty times on Core i7 2.5GHz Linux machine with 8 GB RAM. Despite the large number of states to visit, all runs terminated in less than 27 seconds, and did not violate the desired criteria of correctness with respect to $\texttt{r}*$ invariants.[7] In the

[7] The table with run results is in the appendix, Figure 27.

absence of intentionally removed active wait loops, non-guarded deallocation has lead to **stuck** state in four out of twenty runs.

Can one implement the RCU with weaker order modifiers? To check this hypothesis, we changed release write $[\texttt{rt}]_{rel}$:= (**fst** rtc, loc) in append to a relaxed one, which resulted in 8 out of 10 test runs getting **stuck**, even without deallocation after update in the writer. We then changed only $[\texttt{lhead}]_{rel}$:= a in the writer to a relaxed version, which led to 10 out of 10 test runs ending up in the **stuck** state, as these changes break synchronization between the writer and the reader threads. The same results are observed when changing acquire-reads to *relaxed* ones in traverse. Without the enforced RA-synchronization, the reader threads do not get their viewfronts updated with writes to locations a–d, so attempts to read from them result in **stuck** state, according to READ-UNINIT rule.

As our semantics is implemented in PLT Redex, these synchronization bugs are easy reproduce: if a program gets stuck or delivers unexpected results, one can retrieve the corresponding execution trace via standard Redex machinery.

RCU via consume-reads The original implementation of the RCU used consume-reads instead of acquire-reads. Our version of RCU employs release/acquire synchronization for the following reason. Currently, our semantics does not support *mutable* local variables, as we did not need them for running litmus tests. In their absence, the only possible way to transfer a pointer's value to the next iteration of a loop is to store it in some location, as it is currently done in the **repeat**-loop of traverse. The downside of using a proper memory location instead of a local variable is that this breaks data-dependency chain, which is required by consume-reads for synchronization. There are no fundamental problems preventing us from adding mutable variables to make proper use of consume-reads, and we plan to do it in the future.

7. Related Work

Existing semantics for C11 and their variations The axiomatic C11 semantics by Batty *et al.* [7] has been adapted for establishing soundness of several program logics for relaxed memory [15,27,45, 49]. While some of these adaptations bear a lot of similarity with operational approach [45], all of them are still based on the notion of partial orders between reads and writes. The recent operational semantics for C11 took steps to *incrementalize* Batty *et al.*'s model, constructing the partial orders in a step-wise manner, checking the consistency axioms at every execution step [36]. This model does not follow the program execution order and allows OTA behaviors.

The semantics for *Strong Release-Acquire* (SRA) model by Lahav *et al.* [26] does not use graphs and consistency axioms, relying instead on *message buffers*, reminiscent to our viewfronts in the way they are used for thread synchronization in our approach. However, Lahav *et al.*'s semantics only targets a (strengthened) subset of C11, restricted to release/acquire-synchronization, thus, sidestepping the intricacies of encoding the meaning of relaxed atomics.

Operational semantics for relaxed memory models A low-level operational model for the total store order (TSO) memory model, which is stronger than C11, has been defined by Owens *et al.*, targeting x86-TSO processors [37]. A more complex model is tackled by Sarkar *et al.*, who provided an operational semantics for the POWER architecture [40]. Finally, the most recent work by Flur *et al.* provides an operational model for the ARMv8 architecture [19]. While in this work we are concerned with semantics of a high-level language (*i.e.*, C/C++), investigating compilation schemes to those low-level models with respect to our semantics is our immediate future work.

Related proposals with respect to operational semantics for relaxed memory, inspired by the TSO model, are based on the idea of write-buffers [10, 11, 16, 23], reminiscent to the buffers we use to define postponed operations in our approach. The idea of write buffers and buffer pools fits operational intuition naturally and was used to prove soundness of a program logic [43], but it is not trivial to adapt for C11-style synchronization, especially for reconciling RA-synchronization and relaxed atomics, as we demonstrated in Sections 3 and 4. An alternative approach to define relaxed behavior is to allow the programmer to manipulate with synchronization orders *explicitly* via program-level annotations [13]. This approach provides a highly generic way of modelling custom synchronization patterns at atomic accesses, although, it does not correspond to any specific standard and is not executable. Due to the intentional possibility to use it for modelling very "relaxed" behaviors via arbitrary speculations, the approach allows OTA behaviors. In contrast, our semantics, tailored to allow for modelling all essential concurrent features of C11, prevents OTA by careful treatment of operation buffers, and is executable.

Semantics for relaxed atomics via event structures The OTA executions are considered a serious issue with the C++11 standard [5, 8, 32], in particular, because there is no well-stated and uniform definition of this phenomenon, which is only characterized in the folklore as "values appearing out of nowhere".

At the moment, several proposals provide treatment for relaxed atomics, avoiding the OTA behavior by presenting models for relaxed memory accesses, which are defined in terms of *event structures* [24,38]. These models allow aggressive optimizations including value-range speculations, without introducing classic out-of-thin-air behaviors. Our semantics does not support all of these compiler transformations yet, (*e.g.*, speculative calculation of an arithmetic expression with symbolic values), but they can be added as additional rules without changing the underlying program state.

In contrast with our semantics, the model of Pichon-Pharabod and Sewell [38] is not realistically executable, as it requires for every read operation in a program execution to consider N events, where N is a size of the value domain. That is, for instance, for reading a 32-bit integer it is 2^{32}. Furthermore, at the moment the model [38] does not account for release/acquire-synchronization.

The model by Jeffrey and Riely [24] does not allow for reorderings of independent reads, which makes it too strong to be efficiently implemented on such architectures Power and ARM. The authors suggest a fix for it, which, however, invalidates some other guarantees their initial proposal provides.

Reasoning about Read-Copy-Update RCU structures have been used recently to showcase program logics [27,44], semantic frameworks [26], and program repair/synthesis methods [35] in the context of C11 concurrency. To the best of our knowledge, no other existing approach provides a way of efficiently debugging them by means of re-tracing executions, exhibiting synchronization issues, as we demonstrated in Section 6.3.

8. Conclusion and Future Work

In this work, we presented a family of operational semantics for modelling C/C++ concurrency features. The encoding of C11-style semantics in our framework is based on the two main ideas: *viewfronts* and *operation buffers*, with their various combinations and elaborations allowing to express specific synchronization mechanisms and language aspects from the C11 standard. Our C11 semantics is executable, which we demonstrated by implementing it in PLT Redex and showcasing with a number of examples.

As our future work, we plan to extend the defined formalism for C11 fences [15] and establish formal results relating executions in our semantics to executions in low-level languages via standard compilation schemes [19, 40, 42], thus, proving that our semantics is weak enough to accommodate them. Next, we are going to employ it as a basis for developing a higher-order program logic

for establishing Hoare specifications and program refinement in the C11 model [47], proving the logic's soundness with respect to our semantics, lifted to sets of traces [12] or logical relations [46]. Finally, we plan to use our operational framework for exploring the ideas of efficiently synthesizing synchronization primitives via bounded model checking [35] and partial order reduction [20].

References

[1] C/C++11 mappings to processors. Available from `https://www.cl.cam.ac.uk/~pes20/cpp/cpp0xmappings.html`.

[2] ISO/IEC 14882:2011. Programming language C++, 2011.

[3] ISO/IEC 9899:2011. Programming language C, 2011.

[4] J. Alglave, L. Maranget, and M. Tautschnig. Herding cats: Modelling, simulation, testing, and data mining for weak memory. *ACM Trans. Program. Lang. Syst.*, 36(2):7:1–7:74, 2014.

[5] M. Batty, K. Memarian, K. Nienhuis, J. Pichon-Pharabod, and P. Sewell. The problem of programming language concurrency semantics. In *ESOP*, volume 9032 of *LNCS*, pages 283–307. Springer, 2015.

[6] M. Batty, K. Memarian, S. Owens, S. Sarkar, and P. Sewell. Clarifying and compiling C/C++ concurrency: from C++11 to POWER. In *POPL*, pages 509–520. ACM, 2012.

[7] M. Batty, S. Owens, S. Sarkar, P. Sewell, and T. Weber. Mathematizing C++ concurrency. In *POPL*, pages 55–66. ACM, 2011.

[8] H.-J. Boehm and B. Demsky. Outlawing Ghosts: Avoiding out-of-thin-air results. In *MSPC*, pages 7:1–7:6. ACM, 2014.

[9] R. Bornat, J. Alglave, and M. J. Parkinson. New lace and arsenic: adventures in weak memory with a program logic. *CoRR*, abs/1512.01416, 2015.

[10] G. Boudol and G. Petri. Relaxed memory models: an operational approach. In *POPL*, pages 392–403. ACM, 2009.

[11] G. Boudol, G. Petri, and B. P. Serpette. Relaxed operational semantics of concurrent programming languages. In *EXPRESS/SOS*, volume 89 of *EPTCS*, pages 19–33, 2012.

[12] S. Brookes. A semantics for concurrent separation logic. *Th. Comp. Sci.*, 375(1-3), 2007.

[13] K. Crary and M. J. Sullivan. A Calculus for Relaxed Memory. In *POPL*, pages 623–636. ACM, 2015.

[14] M. Desnoyers, P. E. McKenney, A. S. Stern, M. R. Dagenais, and J. Walpole. User-level implementations of read-copy update. *IEEE Transactions on Parallel and Distributed Systems*, 23(2):375–382, Feb 2012.

[15] M. Doko and V. Vafeiadis. A program logic for C11 memory fences. In *VMCAI*, volume 9583 of *LNCS*, pages 413–430. Springer, 2016.

[16] L. Effinger-Dean and D. Grossman. Modular metatheory for memory consistency models. Technical Report UW-CSE-11-02-01, University of Washington.

[17] M. Felleisen, R. B. Findler, and M. Flatt. *Semantics Engineering with PLT Redex*. MIT Press, 2009.

[18] M. Felleisen and R. Hieb. The revised report on the syntactic theories of sequential control and state. *Theor. Comput. Sci.*, 103(2):235–271, 1992.

[19] S. Flur, K. E. Gray, C. Pulte, S. Sarkar, A. Sezgin, L. Maranget, W. Deacon, and P. Sewell. Modelling the ARMv8 Architecture, Operationally: Concurrency and ISA. In *POPL*, pages 608–621. ACM, 2016.

[20] P. Godefroid. *Partial-Order Methods for the Verification of Concurrent Systems – An Approach to the State-Explosion Problem*. PhD thesis, University of Liege, 1995.

[21] J. L. Hennessy and D. A. Patterson. *Computer Architecture - A Quantitative Approach (5. ed.)*. Morgan Kaufmann, 2012.

[22] C. Hritcu, J. Hughes, B. C. Pierce, A. Spector-Zabusky, D. Vytiniotis, A. Azevedo de Amorim, and L. Lampropoulos. Testing noninterference, quickly. In *ICFP*, pages 455–468. ACM, 2013.

[23] R. Jagadeesan, C. Pitcher, and J. Riely. Generative operational semantics for relaxed memory models. In *ESOP*, volume 6012 of *LNCS*, pages 307–326. Springer, 2010.

[24] A. Jeffrey and J. Riely. On thin air reads: Towards an event structures model of relaxed memory. In *LICS*. IEEE, 2016. To appear.

[25] C. Klein, J. Clements, C. Dimoulas, C. Eastlund, M. Felleisen, M. Flatt, J. A. McCarthy, J. Rafkind, S. Tobin-Hochstadt, and R. B. Findler. Run your research: on the effectiveness of lightweight mechanization. In *POPL*, pages 285–296. ACM, 2012.

[26] O. Lahav, N. Giannarakis, and V. Vafeiadis. Taming Release-Acquire Consistency. In *POPL*, pages 649–662. ACM, 2016.

[27] O. Lahav and V. Vafeiadis. Owicki-Gries Reasoning for Weak Memory Models. In *ICALP (2)*, volume 9135 of *LNCS*, pages 311–323. Springer, 2015.

[28] L. Lamport. How to make a multiprocessor computer that correctly executes multiprocess programs. *IEEE Trans. Computers*, 28(9):690–691, 1979.

[29] L. Maranget, S. Sarkar, and P. Sewell. A Tutorial Introduction to the ARM and POWER Relaxed Memory Models, 2012. Available from `http://www.cl.cam.ac.uk/~pes20/ppc-supplemental/test7.pdf`.

[30] F. Mattern. Virtual time and global states of distributed systems. In *Parallel And Distributed Algorithms*, pages 215–226. North-Holland, 1988.

[31] P. E. McKenney. *Exploiting Deferred Destruction: An Analysis of Read-Copy-Update Techniques in Operating System Kernels*. PhD thesis, OGI School of Science and Engineering at Oregon Health and Sciences University, 2004.

[32] P. E. McKenney, A. Jeffrey, and A. Sezgin. Out-of-thin-air execution is vacuous. Available from `http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2015/n4375.html`.

[33] P. E. McKenney, T. Riegel, J. Preshing, H. Boehm, C. Nelson, and O. Giroux. N4215: Towards implementation and use of `memory_order_consume`, 2014.

[34] P. E. McKenney and J. D. Slingwine. Read-copy update: Using execution history to solve concurrency problems. In *PDCS*, pages 509–518, 1998.

[35] Y. Meshman, N. Rinetzky, and E. Yahav. Pattern-based Synthesis of Synchronization for the C++ Memory Model. In *FMCAD*, pages 120–127, 2015.

[36] K. Nienhuis, K. Memarian, and P. Sewell. An operational semantics for C/C++11 concurrency, 2016. Unpublished draft. Available from `https://www.cl.cam.ac.uk/~pes20/Stuff/c11op.pdf`.

[37] S. Owens, S. Sarkar, and P. Sewell. A Better x86 Memory Model: x86-TSO. In *TPHOLs*, volume 5674 of *LNCS*, pages 391–407. Springer, 2009.

[38] J. Pichon-Pharabod and P. Sewell. A concurrency semantics for relaxed atomics that permits optimisation and avoids thin-air executions. In *POPL*, pages 622–633. ACM, 2016.

[39] E. Pozniansky and A. Schuster. Efficient on-the-fly data race detection in multihreaded C++ programs. In *PPOPP*, pages 179–190. ACM, 2003.

[40] S. Sarkar, P. Sewell, J. Alglave, L. Maranget, and D. Williams. Understanding POWER multiprocessors. In *PLDI*, pages 175–186. ACM, 2011.

[41] I. Sergey, A. Nanevski, and A. Banerjee. Specifying and verifying concurrent algorithms with histories and subjectivity. In *ESOP*, volume 9032 of *LNCS*, pages 333–358. Springer, 2015.

[42] P. Sewell, S. Sarkar, S. Owens, F. Z. Nardelli, and M. O. Myreen. x86-tso: a rigorous and usable programmer's model for x86 multiprocessors. *Commun. ACM*, 53(7):89–97, 2010.

[43] F. Sieczkowski, K. Svendsen, L. Birkedal, and J. Pichon-Pharabod. A separation logic for fictional sequential consistency. In *ESOP*, volume 9032 of *LNCS*, pages 736–761. Springer, 2015.

[44] J. Tassarotti, D. Dreyer, and V. Vafeiadis. Verifying Read-Copy-Update in a Logic for Weak Memory. In *PLDI*, pages 110–120. ACM, 2015.

[45] A. Turon, V. Vafeiadis, and D. Dreyer. GPS: navigating weak memory with ghosts, protocols, and separation. In *OOPSLA*, pages 691–707. ACM, 2014.

[46] A. J. Turon, J. Thamsborg, A. Ahmed, L. Birkedal, and D. Dreyer. Logical relations for fine-grained concurrency. In *POPL*, pages 343–356. ACM, 2013.

[47] V. Vafeiadis. Formal reasoning about the C11 weak memory model. In *CPP*, pages 1–2. ACM, 2015.

[48] V. Vafeiadis, T. Balabonski, S. Chakraborty, R. Morisset, and F. Z. Nardelli. Common Compiler Optimisations are Invalid in the C11 Memory Model and what we can do about it. In *POPL*, pages 209–220. ACM, 2015.

[49] V. Vafeiadis and C. Narayan. Relaxed Separation Logic: a program logic for C11 concurrency. In *OOPSLA*, pages 867–884. ACM, 2013.
[50] D. L. Weaver and T. Germond. *The SPARC Architecture Manual*. SPARC International, Inc., 1994. Version 9.

A. The Catalogue of Litmus Tests

A.1 Store Buffering (SB)

SB_rel+acq
Fully Supported: ✓
Requires: History + Viewfronts

Possible outcomes:
`r1 = 0 ∧ r2 = 0`
`r1 = 0 ∧ r2 = 1`
`r1 = 1 ∧ r2 = 0`
`r1 = 1 ∧ r2 = 1`

$[\mathtt{x}]_{rel} := 0;\ [\mathtt{y}]_{rel} := 0;$
$[\mathtt{x}]_{rel} := 1;\ \|\ [\mathtt{y}]_{rel} := 1;$
$\mathtt{r1} = [\mathtt{y}]_{acq}\ \|\ \mathtt{r2} = [\mathtt{x}]_{acq}$

SB_sc
Fully Supported: ✓
Requires: SC + History + Viewfronts

Forbidden outcomes:
`r1 = 0 ∧ r2 = 0`

$[\mathtt{x}]_{sc} := 0;\ [\mathtt{y}]_{sc} := 0;$
$[\mathtt{x}]_{sc} := 1;\ \|\ [\mathtt{y}]_{sc} := 1;$
$\mathtt{r1} = [\mathtt{y}]_{sc}\ \|\ \mathtt{r2} = [\mathtt{x}]_{sc}$

SB_sc+rel
Fully Supported: ✓
Requires: SC + History + Viewfronts

Possible outcomes:
`r1 = 0 ∧ r2 = 0`
`r1 = 0 ∧ r2 = 1`
`r1 = 1 ∧ r2 = 0`
`r1 = 1 ∧ r2 = 1`

$[\mathtt{x}]_{sc} := 0;\ [\mathtt{y}]_{sc} := 0;$
$[\mathtt{x}]_{rel} := 1;\ \|\ [\mathtt{y}]_{sc} := 1;$
$\mathtt{r1} = [\mathtt{y}]_{sc}\ \|\ \mathtt{r2} = [\mathtt{x}]_{sc}$

SB_sc+acq
Fully Supported: ✓
Requires: SC + History + Viewfronts

Possible outcomes:
`r1 = 0 ∧ r2 = 0`
`r1 = 0 ∧ r2 = 1`
`r1 = 1 ∧ r2 = 0`
`r1 = 1 ∧ r2 = 1`

$[\mathtt{x}]_{sc} := 0;\ [\mathtt{y}]_{sc} := 0;$
$[\mathtt{x}]_{sc} := 1;\ \|\ [\mathtt{y}]_{sc} := 1;$
$\mathtt{r1} = [\mathtt{y}]_{acq}\ \|\ \mathtt{r2} = [\mathtt{x}]_{sc}$

A.2 Load Buffering (LB)

LB_rlx
Fully Supported: ✓
Requires: Postponed Reads + History + Viewfronts

Possible outcomes:
`r1 = 0 ∧ r2 = 0`
`r1 = 0 ∧ r2 = 1`
`r1 = 1 ∧ r2 = 0`
`r1 = 1 ∧ r2 = 1`

$[\mathtt{x}]_{rlx} := 0;\ [\mathtt{y}]_{rlx} := 0;$
$\mathtt{r1} = [\mathtt{y}]_{rlx};\ \|\ \mathtt{r2} = [\mathtt{x}]_{rlx};$
$[\mathtt{x}]_{rlx} := 1\ \|\ [\mathtt{y}]_{rlx} := 1$

LB_rel+rlx
Fully Supported: ✓
Requires: Postponed Reads + History + Viewfronts

Possible outcomes:
`r1 = 0 ∧ r2 = 0`
`r1 = 0 ∧ r2 = 1`
`r1 = 1 ∧ r2 = 0`
`r1 = 1 ∧ r2 = 1`

$[\mathtt{x}]_{rlx} := 0;\ [\mathtt{y}]_{rlx} := 0;$
$\mathtt{r1} = [\mathtt{y}]_{rlx};\ \|\ \mathtt{r2} = [\mathtt{x}]_{rlx};$
$[\mathtt{x}]_{rel} := 1\ \|\ [\mathtt{y}]_{rel} := 1$

LB_rel+rlx
Fully Supported: ✓
Requires: Postponed Reads + History + Viewfronts

Possible outcomes:
`r1 = 0 ∧ r2 = 0`
`r1 = 0 ∧ r2 = 1`
`r1 = 1 ∧ r2 = 0`
`r1 = 1 ∧ r2 = 1`

$[\mathtt{x}]_{rlx} := 0;\ [\mathtt{y}]_{rlx} := 0;$
$\mathtt{r1} = [\mathtt{y}]_{rlx};\ \|\ \mathtt{r2} = [\mathtt{x}]_{rlx};$
$[\mathtt{x}]_{rel} := 1\ \|\ [\mathtt{y}]_{rel} := 1$

LB_acq+rlx
Fully Supported: ✗
Requires: Postponed Reads + History + Viewfronts

Possible outcomes:
`r1 = 0 ∧ r2 = 0`
`r1 = 0 ∧ r2 = 1`
`r1 = 1 ∧ r2 = 0`
`r1 = 1 ∧ r2 = 1`

$[\mathtt{x}]_{rlx} := 0;\ [\mathtt{y}]_{rlx} := 0;$
$\mathtt{r1} = [\mathtt{y}]_{acq};\ \|\ \mathtt{r2} = [\mathtt{x}]_{acq};$
$[\mathtt{x}]_{rlx} := 1\ \|\ [\mathtt{y}]_{rlx} := 1$

Our semantics doesn't allow the `r1 = 1 ∧ r2 = 1` outcome for the program. It doesn't allow reordering of an acquire read with a subsequent write. The known sound compilation schemes of acquire read to major platforms (x86, ARM, Power) don't allow the behavior either.

LB_rel+acq+rlx
Fully Supported: ✓
Requires: Postponed Reads + History + Viewfronts

Forbidden outcomes:
`r1 = 1 ∧ r2 = 1`

$[\mathtt{x}]_{rlx} := 0;\ [\mathtt{y}]_{rlx} := 0;$
$\mathtt{r1} = [\mathtt{y}]_{acq};\ \|\ \mathtt{r2} = [\mathtt{x}]_{rlx};$
$[\mathtt{x}]_{rlx} := 1\ \|\ [\mathtt{y}]_{rel} := 1$

LB_rlx+use
Fully Supported: ✓
Requires: Postponed Reads + History + Viewfronts

Allowed outcome:
`r1 = 1 ∧ r2 = 1`

$[\mathtt{x}]_{rlx} := 0;\ [\mathtt{y}]_{rlx} := 0;$

$\mathtt{r1} = [\mathtt{y}]_{rlx};\ [\mathtt{z1}]_{rlx} := \mathtt{r1};\ [\mathtt{x}]_{rlx} := 1 \ \|\ \mathtt{r2} = [\mathtt{x}]_{rlx};\ [\mathtt{z2}]_{rlx} := \mathtt{r2};\ [\mathtt{y}]_{rlx} := 1$

LB_rlx+let
Fully Supported: ✓
Requires: Postponed Reads + History + Viewfronts

Allowed outcome:
`r1 = 1 ∧ r'1 = 2 ∧ r2 = 1 ∧ r'2 = 2`

$[\mathtt{x}]_{rlx} := 0;\ [\mathtt{y}]_{rlx} := 0;$

$\mathtt{r1} = [\mathtt{y}]_{rlx};\ \mathtt{r'1} = \mathtt{r1} + 1;\ [\mathtt{x}]_{rlx} := 1 \ \|\ \mathtt{r2} = [\mathtt{x}]_{rlx};\ \mathtt{r'2} = \mathtt{r2} + 1;\ [\mathtt{y}]_{rlx} := 1$

LB_rlx+join
Fully Supported: ✓
Requires: Postponed Reads + History + Viewfronts + JN

Allowed outcomes:
`r1 = 1 ∧ r2 = 1`

$[\mathtt{x}]_{rlx} := 0;\ [\mathtt{y}]_{rlx} := 0;$

$\mathtt{r1} = [\mathtt{y}]_{rlx};\ [\mathtt{z1}]_{rlx} := \mathtt{r1} \ \|\ 0 \ \|\ [\mathtt{x}]_{rlx} := 1 \ \|\ \mathtt{r2} = [\mathtt{x}]_{rlx};\ [\mathtt{z2}]_{rlx} := \mathtt{r2} \ \|\ 0 \ \|\ [\mathtt{y}]_{rlx} := 1$

LB_rel+rlx+join
Fully Supported: ✓
Requires: Postponed Reads + History + Viewfronts + JN

Allowed outcomes:
`r1 = 1 ∧ r2 = 1`

$[\mathtt{x}]_{rlx} := 0;\ [\mathtt{y}]_{rlx} := 0;$

$\mathtt{r1} = [\mathtt{y}]_{rlx};\ [\mathtt{z1}]_{rlx} := \mathtt{r1} \ \|\ 0 \ \|\ [\mathtt{x}]_{rel} := 1 \ \|\ \mathtt{r2} = [\mathtt{x}]_{rlx};\ [\mathtt{z2}]_{rlx} := \mathtt{r2} \ \|\ 0 \ \|\ [\mathtt{y}]_{rel} := 1$

LB_acq+rlx+join
Fully Supported: ✗
Requires: Postponed Reads + History + Viewfronts + JN

Allowed outcomes:
`r1 = 1 ∧ r2 = 1`

$[\mathtt{x}]_{rlx} := 0;\ [\mathtt{y}]_{rlx} := 0;$

$\mathtt{r1} = [\mathtt{y}]_{acq};\ [\mathtt{z1}]_{rlx} := \mathtt{r1} \ \|\ 0 \ \|\ [\mathtt{x}]_{rlx} := 1 \ \|\ \mathtt{r2} = [\mathtt{x}]_{acq};\ [\mathtt{z2}]_{rlx} := \mathtt{r2} \ \|\ 0 \ \|\ [\mathtt{y}]_{rlx} := 1$

A.3 Message Passing (MP)

MP_rlx+na
Fully Supported: ✓
Requires: NA + History + Viewfronts

Possible outcomes:
`r1 = 0`
`r1 = 5`
`stuck`

$[\mathtt{f}]_{rlx} := 0;\ [\mathtt{d}]_{na} := 0;$

$[\mathtt{d}]_{na} := 5;\ [\mathtt{f}]_{rlx} := 1 \ \|\ \mathbf{repeat}\ [\mathtt{f}]_{rlx}\ \mathbf{end};\ \mathtt{r1} = [\mathtt{d}]_{na}$

MP_rel+rlx+na
Fully Supported: ✓
Requires: NA + History + Viewfronts

Possible outcomes:
`r1 = 0`
`r1 = 5`
`stuck`

$[\mathtt{f}]_{rlx} := 0;\ [\mathtt{d}]_{na} := 0;$

$[\mathtt{d}]_{na} := 5;\ [\mathtt{f}]_{rel} := 1 \ \|\ \mathbf{repeat}\ [\mathtt{f}]_{rlx}\ \mathbf{end};\ \mathtt{r1} = [\mathtt{d}]_{na}$

MP_rlx+acq+na
Fully Supported: ✓
Requires: NA + History + Viewfronts

Possible outcomes:
`r1 = 0`
`r1 = 5`
`stuck`

$[\mathtt{f}]_{rlx} := 0;\ [\mathtt{d}]_{na} := 0;$

$[\mathtt{d}]_{na} := 5;\ [\mathtt{f}]_{rlx} := 1 \ \|\ \mathbf{repeat}\ [\mathtt{f}]_{acq}\ \mathbf{end};\ \mathtt{r1} = [\mathtt{d}]_{na}$

MP_rel+acq+na
Fully Supported: ✓
Requires: NA + History + Viewfronts

Possible outcomes:
`r1 = 5`

$[\mathtt{f}]_{rel} := 0;\ [\mathtt{d}]_{na} := 0;$

$[\mathtt{d}]_{na} := 5;\ [\mathtt{f}]_{rel} := 1 \ \|\ \mathbf{repeat}\ [\mathtt{f}]_{acq}\ \mathbf{end};\ \mathtt{r1} = [\mathtt{d}]_{na}$

MP_rel+acq+na+rlx
Fully Supported: ✓
Requires: Write-fronts + NA + History + Viewfronts

Possible outcomes:
`r1 = 5`

$[\mathtt{f}]_{rel} := 0;\ [\mathtt{d}]_{na} := 0;$

$[\mathtt{d}]_{na} := 5;\ [\mathtt{f}]_{rel} := 1;\ [\mathtt{f}]_{rlx} := 2 \ \|\ \mathbf{repeat}\ [\mathtt{f}]_{acq} == 2\ \mathbf{end};\ \mathtt{r1} = [\mathtt{d}]_{na}$

MP_rel+acq+na+rlx_2
Fully Supported: ✓
Requires: Write-fronts + NA + History + Viewfronts

Possible outcomes:
`r1 = 5 ∧ r2 = <0, 1>`

$[\mathtt{f}]_{na} := 0;\ [\mathtt{d}]_{na} := 0;\ [\mathtt{x}]_{na} := 0;$

$[\mathtt{d}]_{na} := 5;\ [\mathtt{f}]_{rel} := 1;\ [\mathtt{x}]_{rel} := 1;\ [\mathtt{f}]_{rlx} := 2$ ‖ **repeat** $[\mathtt{f}]_{acq} == 2$ **end**; r1 := $[\mathtt{d}]_{na}$; r2 := $[\mathtt{x}]_{rlx}$

MP_con+na
Fully Supported: ✓
Requires: Consume + NA + History + Viewfronts

Possible outcomes:
`r1 = 0`
`r1 = 5`

$[\mathtt{f}]_{con}$:= **null**; $[\mathtt{d}]_{na} := 0;$

$[\mathtt{d}]_{na} := 5;\ [\mathtt{f}]_{rel} := \mathtt{d}$ ‖ r0 := $[\mathtt{f}]_{con}$; **if** r0 != **null** **then** r1 = $[\mathtt{r0}]_{na}$ **else** r1 = 0 **fi**

MP_con+na_2
Fully Supported: ✓
Requires: Consume + NA + History + Viewfronts

Possible outcomes:
`r2 = 0 ∧ r3 = <0, 1>`
`r2 = 5 ∧ r3 = <0, 1>`

$[\mathtt{p}]_{na}$:= **null**; $[\mathtt{d}]_{na} := 0;\ [\mathtt{x}]_{na} := 0;$

$[\mathtt{x}]_{rlx} := 1;\ [\mathtt{d}]_{na} := 1;\ [\mathtt{p}]_{rel} := \mathtt{d}$ ‖ r1 = $[\mathtt{p}]_{con}$; **if** r1 != **null** **then** r2 = $[\mathtt{r1}]_{na}$; r3 = $[\mathtt{x}]_{rlx}$ **else** r2 = 0; r3 = 0 **fi**

MP_cas+rel+acq+na from [49]
Fully Supported: ✓
Requires: NA + History + Viewfronts

Impossible outcomes:
stuck

$[\mathtt{f}]_{rlx} := 1;\ [\mathtt{d}]_{na} := 0;$

$[\mathtt{d}]_{na} := 5;\ [\mathtt{f}]_{rel} := 0$ ‖ r1 = $\mathtt{cas}_{acq,rlx}(\mathtt{f}, 0, 1)$; **if** r1 == 0 **then** $[\mathtt{d}]_{rlx} := 6$ **else** 0 **fi** ‖ r2 = $\mathtt{cas}_{acq,rlx}(\mathtt{f}, 0, 1)$; **if** r2 == 0 **then** $[\mathtt{d}]_{rlx} := 7$ **else** 0 **fi**

MP_cas+rel+rlx+na
Fully Supported: ✓
Requires: NA + History + Viewfronts

Possible outcomes:
stuck

$[\mathtt{f}]_{rlx} := 1;\ [\mathtt{d}]_{na} := 0;$

$[\mathtt{d}]_{na} := 5\ [\mathtt{f}]_{rel} := 0;$ ‖ r1 = $\mathtt{cas}_{rlx,rlx}(\mathtt{f}, 0, 1)$; **if** r1 == 0 **then** $[\mathtt{d}]_{rlx} := 6$ **else** 0 **fi** ‖ r2 = $\mathtt{cas}_{rlx,rlx}(\mathtt{f}, 0, 1)$; **if** r2 == 0 **then** $[\mathtt{d}]_{rlx} := 7$ **else** 0 **fi**

A.4 Coherence of Read-Read (CoRR)

CoRR_rlx
Fully Supported: ✓
Requires: History + Viewfronts

Impossible outcomes:
`r1 = 1 ∧ r2 = 2 ∧ r3 = 2 ∧ r4 = 1`
`r1 = 2 ∧ r2 = 1 ∧ r3 = 1 ∧ r4 = 2`

$[\mathtt{x}]_{rlx} := 0;$

$[\mathtt{x}]_{rlx} := 1$ ‖ $[\mathtt{x}]_{rlx} := 2$ ‖ r1 = $[\mathtt{x}]_{rlx}$; r2 = $[\mathtt{x}]_{rlx}$ ‖ r3 = $[\mathtt{x}]_{rlx}$; r4 = $[\mathtt{x}]_{rlx}$

CoRR_rel+acq
Fully Supported: ✓
Requires: History + Viewfronts

Impossible outcomes:
`r1 = 1 ∧ r2 = 2 ∧ r3 = 2 ∧ r4 = 1`
`r1 = 2 ∧ r2 = 1 ∧ r3 = 1 ∧ r4 = 2`

$[\mathtt{x}]_{rel} := 0;$

$[\mathtt{x}]_{rel} := 1$ ‖ $[\mathtt{x}]_{rel} := 2$ ‖ r1 = $[\mathtt{x}]_{acq}$; r2 = $[\mathtt{x}]_{acq}$ ‖ r3 = $[\mathtt{x}]_{acq}$; r4 = $[\mathtt{x}]_{acq}$

A.5 Independent Reads of Independent Writes (IRIW)

IRIW_rlx
Fully Supported: ✓
Requires: History + Viewfronts

Possible outcomes:
`r1 = <0, 1>; r2 = <0, 1>;`
`r3 = <0, 1>; r4 = <0, 1>`

$[\mathtt{x}]_{rlx} := 0;\ [\mathtt{y}]_{rlx} := 0;$

$[\mathtt{x}]_{rlx} := 1$ ‖ $[\mathtt{y}]_{rlx} := 1$ ‖ r1 = $[\mathtt{x}]_{rlx}$; r2 = $[\mathtt{y}]_{rlx}$ ‖ r3 = $[\mathtt{y}]_{rlx}$; r4 = $[\mathtt{x}]_{rlx}$

Comment: It is possible to get `r1 = 1; r2 = 0; r3 = 1; r4 = 0`

IRIW_rel+acq
Fully Supported: ✓
Requires: History + Viewfronts

Possible outcomes:
`r1 = <0, 1>; r2 = <0, 1>;`
`r3 = <0, 1>; r4 = <0, 1>`

$[\mathtt{x}]_{rel} := 0;\ [\mathtt{y}]_{rel} := 0;$

$[\mathtt{x}]_{rel} := 1$ ‖ $[\mathtt{y}]_{rel} := 1$ ‖ r1 = $[\mathtt{x}]_{acq}$; r2 = $[\mathtt{y}]_{acq}$ ‖ r3 = $[\mathtt{y}]_{acq}$; r4 = $[\mathtt{x}]_{acq}$

Comment: It is possible to get `r1 = 1; r2 = 0; r3 = 1; r4 = 0`

IRIW_sc
Fully Supported: ✓
Requires: SC + History + Viewfronts

Forbidden outcomes:
`r1 = 1 ∧ r2 = 0 ∧ r3 = 1 ∧ r4 = 0`

$$\begin{array}{l} [\mathtt{x}]_{sc} := 0;\ [\mathtt{y}]_{sc} := 0; \\ [\mathtt{x}]_{sc} := 1 \ \Big\|\ [\mathtt{y}]_{sc} := 1 \ \Big\|\ \begin{array}{l}\mathtt{r1} = [\mathtt{x}]_{sc};\\ \mathtt{r2} = [\mathtt{y}]_{sc}\end{array} \ \Big\|\ \begin{array}{l}\mathtt{r3} = [\mathtt{y}]_{sc};\\ \mathtt{r4} = [\mathtt{x}]_{sc}\end{array} \end{array}$$

A.6 Write-to-Read Causality (WRC)

WRC_rel+acq
Fully Supported: ✓
Requires: History + Viewfronts

Forbidden outcomes:
`r2 = 1 ∧ r3 = 0`

$$\begin{array}{l} [\mathtt{x}]_{rel} := 0;\ [\mathtt{y}]_{rel} := 0; \\ [\mathtt{x}]_{rel} := 1 \ \Big\|\ \begin{array}{l}\mathtt{r1} = [\mathtt{x}]_{acq};\\ [\mathtt{y}]_{rel} := \mathtt{r1}\end{array} \ \Big\|\ \begin{array}{l}\mathtt{r2} = [\mathtt{y}]_{acq};\\ \mathtt{r3} = [\mathtt{x}]_{acq}\end{array} \end{array}$$

WRC_rlx
Fully Supported: ✓
Requires: History + Viewfronts

Possible outcomes:
`r2 = 0 ∧ r3 = 0`
`r2 = 0 ∧ r3 = 1`
`r2 = 1 ∧ r3 = 0`
`r2 = 1 ∧ r3 = 1`

$$\begin{array}{l} [\mathtt{x}]_{rlx} := 0;\ [\mathtt{y}]_{rlx} := 0; \\ [\mathtt{x}]_{rlx} := 1 \ \Big\|\ \begin{array}{l}\mathtt{r1} = [\mathtt{x}]_{rlx};\\ [\mathtt{y}]_{rlx} := \mathtt{r1}\end{array} \ \Big\|\ \begin{array}{l}\mathtt{r2} = [\mathtt{y}]_{rlx};\\ \mathtt{r3} = [\mathtt{x}]_{rlx}\end{array} \end{array}$$

`r2 = 1; r3 = 0`

WRC_cas+rel
Fully Supported: ✓
Requires: History + Viewfronts

Impossible outcomes:
`r2 = 2 ∧ r3 = 0`

$$\begin{array}{l} [\mathtt{x}]_{rel} := 0;\ [\mathtt{y}]_{rel} := 0; \\ \begin{array}{l}[\mathtt{x}]_{rel} := 1;\\ [\mathtt{y}]_{rel} := 1\end{array} \ \Big\|\ \mathtt{cas}_{rel,acq}(\mathtt{y},\ 1,\ 2) \ \Big\|\ \begin{array}{l}\mathtt{r1} = [\mathtt{y}]_{rel};\\ \mathtt{r2} = [\mathtt{x}]_{rel}\end{array} \end{array}$$

WRC_cas+rlx
Fully Supported: ✓
Requires: History + Viewfronts

Impossible outcomes:
`r2 = 2 ∧ r3 = 0`

$$\begin{array}{l} [\mathtt{x}]_{rlx} := 0;\ [\mathtt{y}]_{rlx} := 0; \\ \begin{array}{l}[\mathtt{x}]_{rlx} := 1;\\ [\mathtt{y}]_{rel} := 1\end{array} \ \Big\|\ \mathtt{cas}_{rlx,rlx}(\mathtt{y},\ 1,\ 2) \ \Big\|\ \begin{array}{l}\mathtt{r1} = [\mathtt{y}]_{rlx};\\ \mathtt{r2} = [\mathtt{x}]_{rlx}\end{array} \end{array}$$

A.7 Out-of-Thin-Air reads

In our semantics it is not possible to get out-of-thin-air results, unlike the C11 standard. But such reads are considered to be an undesirable behavior by most of the standard's clients [5].

OTA_lb
Fully Supported: ✗
Requires: Postponed reads + History + Viewfronts

Possible outcomes:
`r1 = 0 ∧ r2 = 0`

$$\begin{array}{l} [\mathtt{x}]_{rlx} := 0;\ [\mathtt{y}]_{rlx} := 0; \\ \begin{array}{l}\mathtt{r1} = [\mathtt{y}]_{rlx};\\ [\mathtt{x}]_{rlx} := \mathtt{r1}\end{array} \ \Big\|\ \begin{array}{l}\mathtt{r2} = [\mathtt{x}]_{rlx};\\ [\mathtt{y}]_{rlx} := \mathtt{r2}\end{array} \end{array}$$

Comment: According to the C11 standard [2, 3], `r1` and `r2` can get arbitrary values.

OTA_if
Fully Supported: ✗
Requires: Postponed reads + History + Viewfronts

Possible outcomes:
`r1 = 0 ∧ r2 = 0`

$$\begin{array}{l} [\mathtt{x}]_{rlx} := 0;\ [\mathtt{y}]_{rlx} := 0; \\ \begin{array}{l}\mathtt{r1} = [\mathtt{y}]_{rlx};\\ \mathbf{if}\ \mathtt{r1}\\ \mathbf{then}\ [\mathtt{x}]_{rlx} := 1\\ \mathbf{else}\ \mathtt{r1} = 0\\ \mathbf{fi}\end{array} \ \Big\|\ \begin{array}{l}\mathtt{r2} = [\mathtt{x}]_{rlx};\\ \mathbf{if}\ \mathtt{r2}\\ \mathbf{then}\ [\mathtt{y}]_{rlx} := 1\\ \mathbf{else}\ \mathtt{r2} = 0\\ \mathbf{fi}\end{array} \end{array}$$

Comment: According to the C11 standard [2, 3], `r1` and `r2` can be 1s at the end of execution.

A.8 Write Reorder (WR), or 2+2W from [26]

WR_rlx
Fully Supported: ✓
Requires: History + Viewfronts + Operational Buffers

Possible outcomes:
`r1 = 1 ∧ r2 = 2`
`r1 = 2 ∧ r2 = 1`
`r1 = 2 ∧ r2 = 2`

$$\begin{array}{l} [\mathtt{x}]_{rlx} := 0;\ [\mathtt{y}]_{rlx} := 0; \\ \begin{array}{l}[\mathtt{x}]_{rlx} := 1;\\ [\mathtt{y}]_{rlx} := 2\end{array} \ \Big\|\ \begin{array}{l}[\mathtt{y}]_{rlx} := 1;\\ [\mathtt{x}]_{rlx} := 2\end{array} \\ \mathtt{r1} = [\mathtt{x}]_{rlx};\ \mathtt{r2} = [\mathtt{y}]_{rlx} \end{array}$$

WR_rlx+rel
Fully Supported: ✓
Requires: History + Viewfronts + Operational Buffers

Possible outcomes:
`r1 = 1 ∧ r2 = 2`
`r1 = 2 ∧ r2 = 1`
`r1 = 2 ∧ r2 = 2`

$[x]_{rlx} := 0;\ [y]_{rlx} := 0;$

$[x]_{rlx} := 1;$	‖	$[y]_{rlx} := 1;$
$[y]_{rel} := 2$		$[x]_{rel} := 2$

$r1 = [x]_{rlx};\ r2 = [y]_{rlx}$

WR_rel
Fully Supported: ✓
Requires: History + Viewfronts + Operational Buffers

Possible outcomes:
`r1 = 1 ∧ r2 = 2`
`r1 = 2 ∧ r2 = 1`
`r1 = 2 ∧ r2 = 2`

$[x]_{rel} := 0;\ [y]_{rel} := 0;$

$[x]_{rel} := 1;$	‖	$[y]_{rel} := 1;$
$[y]_{rel} := 2$		$[x]_{rel} := 2$

$r1 = [x]_{acq};\ r2 = [y]_{acq}$

A.9 Speculative Execution

SE_simple
Fully Supported: ✓
Requires:

Possible outcomes:
`r0 = 0`
`r0 = 1`

$[x]_{rlx} := 0;\ [y]_{rlx} := 0;\ [z]_{rlx} := 0;$

r1 = $[x]_{rlx}$;	‖	r2 = $[y]_{rlx}$;
if r1		**if** r2
then $[z]_{rlx} := 1;$		**then** $[x]_{rlx} := 1$
$[y]_{rlx} := 1$		**else** 0
else $[y]_{rlx} := 1$		**fi**
fi		

$r0 = [z]_{rlx}$

SE_prop
Fully Supported: ✓
Requires:

Possible outcomes:
`r0 = 0`
`r0 = 1`

$[x]_{rlx} := 0;\ [y]_{rlx} := 0;\ [z]_{rlx} := 0;$

r1 = $[x]_{rlx}$;	‖	
if r1		r2 = $[y]_{rlx}$;
then $[z]_{rlx} := 1;$		**if** r2
r1 = $[z]_{rlx}$;		**then** $[x]_{rlx} := 1$
$[y]_{rlx} := r1$		**else** 0
else $[y]_{rlx} := 1$		**fi**
fi		

$r0 = [z]_{rlx}$

SE_nested
Fully Supported: ✓
Requires:

Possible outcomes:
`r0 = 0`
`r0 = 1`

$[x]_{rlx} := 0;\ [y]_{rlx} := 0;\ [z]_{rlx} := 0;\ [f]_{rlx} := 0;$

r1 = $[x]_{rlx}$;	‖	
if r1		
then r2 = $[f]_{rlx}$;		r3 = $[y]_{rlx}$;
if r2		**if** r3
then $[z]_{rlx} := 1;$		**then** $[f]_{rlx} := 1;$
$[y]_{rlx} := 1$		$[x]_{rlx} := 1$
else $[y]_{rlx} := 1$		**else** 0
fi		**fi**
else $[y]_{rlx} := 1$		
fi		

$r0 = [z]_{rlx}$

A.10 Locks

Dekker's lock

Possible outcomes:
stuck
Requires: RA + na
Fully Supported: ✓

$[x]_{rel} := 0;\ [y]_{rel} := 0;\ [d]_{na} := 0;$

$[x]_{rel} := 1;$	‖	$[y]_{rel} := 1;$
r1 = $[y]_{acq}$		r2 = $[x]_{acq}$
if r1 == 0		**if** r2 == 0
then $[d]_{na} := 5$		**then** $[d]_{na} := 6$
else 0		**else** 0
fi		**fi**

Cohen's lock

Impossible outcomes (according to [45]):
  stuck
Requires: RA + na
Fully Supported: ✓

```
[x]_rel := 0; [y]_rel := 0; [d]_na := 0;
```

[x]$_{rel}$:= **choice** 1 2;	[y]$_{rel}$:= **choice** 1 2;
repeat [y]$_{acq}$ **end**;	**repeat** [x]$_{acq}$ **end**;
r1 = [x]$_{acq}$	r3 = [x]$_{acq}$
r2 = [y]$_{acq}$	r4 = [y]$_{acq}$
if r1 == r2	**if** r3 != r4
then [d]$_{na}$:= 5	**then** [d]$_{na}$:= 6
else 0	**else** 0
fi	**fi**

B. Additional Semantic Rules

$$
\begin{array}{rcl}
\mathsf{E} & ::= & [\,] \mid \mathsf{x} = \mathsf{E}; \mathsf{s} \\
 & & \mid \textbf{par}\ \mathsf{E}\ \mathsf{s} \mid \textbf{par}\ \mathsf{s}\ \mathsf{E} \\
\mathsf{EU} & ::= & [\,] \mid (\mathsf{EU}) \mid \mathsf{EU}\ \mathsf{op}\ \mathsf{e} \mid \mathsf{e}\ \mathsf{op}\ \mathsf{EU} \\
 & & \mid (\mu,\ \mathsf{EU}) \mid (\mathsf{EU},\ \mu) \\
 & & \mid \textbf{fst}\ \mathsf{EU} \mid \textbf{snd}\ \mathsf{EU} \\
 & & \mid \textbf{choice}\ \mathsf{EU}\ \mathsf{e} \mid \textbf{choice}\ \mathsf{e}\ \mathsf{EU} \\
 & & \mid \mathsf{x} = \mathsf{EU}; \mathsf{s} \\
 & & \mid \textbf{if}\ \mathsf{EU}\ \textbf{then}\ \mathsf{s}_1\ \textbf{else}\ \mathsf{s}_2\ \textbf{fi} \\
 & & \mid [\iota]_{\mathsf{WM}} := \mathsf{EU} \\
 & & \mid \mathsf{cas}_{\mathsf{SM},\mathsf{FM}}(\iota, \mathsf{EU}, \mu) \mid \mathsf{cas}_{\mathsf{SM},\mathsf{FM}}(\iota, \mu, \mathsf{EU})
\end{array}
$$

Figure 23. Syntax of evaluation contexts.

$$\frac{}{\langle \mathsf{E}[\mathsf{x} = v; \mathsf{s}], \xi\rangle \Longrightarrow \langle \mathsf{E}[\mathsf{s}[\mathsf{x}/v]], \xi\rangle}\ \textsc{Subst}$$

$$\frac{n \neq 0}{\langle \mathsf{E}[\textbf{if}\ n\ \textbf{then}\ \mathsf{s}_1\ \textbf{else}\ \mathsf{s}_2\ \textbf{fi}], \xi\rangle \Longrightarrow \langle \mathsf{E}[\mathsf{s}_1]], \xi\rangle}\ \textsc{If-True}$$

$$\frac{}{\langle \mathsf{E}[\textbf{if}\ 0\ \textbf{then}\ \mathsf{s}_1\ \textbf{else}\ \mathsf{s}_2\ \textbf{fi}], \xi\rangle \Longrightarrow \langle \mathsf{E}[\mathsf{s}_2]], \xi\rangle}\ \textsc{If-False}$$

$$\frac{\mathsf{x} \text{ – fresh invariable}}{\langle \mathsf{E}[\textbf{repeat}\ \mathsf{s}\ \textbf{end}], \xi\rangle \Longrightarrow \langle \mathsf{E}[\mathsf{x} = \mathsf{s}; \textbf{if}\ \mathsf{x}\ \textbf{then}\ \mathsf{x}\ \textbf{else repeat}\ \mathsf{s}\ \textbf{end fi}], \xi\rangle}\ \textsc{Repeat-Unroll}$$

$$\frac{\xi' = \mathsf{spawn}(\mathsf{E},\ \xi)}{\langle \mathsf{E}[\textbf{spw}\ \mathsf{s}_1\ \mathsf{s}_2], \xi\rangle \Longrightarrow \langle \mathsf{E}[\textbf{par}\ \mathsf{s}_1\ \mathsf{s}_2], \xi'\rangle}\ \textsc{Spawn} \qquad \frac{\xi' = \mathsf{join}(\mathsf{E},\ \xi)}{\langle \mathsf{E}[\textbf{par}\ v_1\ v_2], \xi\rangle \Longrightarrow \langle \mathsf{E}[(v_1,\ v_2)], \xi'\rangle}\ \textsc{Join}$$

$$\frac{}{\langle \mathsf{E}[\mathsf{EU}[\textbf{choice}\ \mathsf{e}_1\ \mathsf{e}_2]], \xi\rangle \Longrightarrow \langle \mathsf{E}[\mathsf{EU}[\mathsf{e}_1]], \xi\rangle}\ \textsc{Choice-Fst} \qquad \frac{}{\langle \mathsf{E}[\mathsf{EU}[\textbf{choice}\ \mathsf{e}_1\ \mathsf{e}_2]], \xi\rangle \Longrightarrow \langle \mathsf{E}[\mathsf{EU}[\mathsf{e}_2]], \xi\rangle}\ \textsc{Choice-Snd}$$

Figure 24. The core rules of the semantics.

$$\frac{\xi = \langle H, \psi^{\mathtt{rd}}, \ldots, \sigma^{\mathtt{na}} \rangle \quad \ldots \quad \sigma_{\mathtt{rd}}(\ell) \equiv \mathsf{LastTS}(H, \ell) \quad \sigma = \sigma_{\mathtt{rd}}[\ell \mapsto \tau] \quad \xi' = \langle H[(\ell, \tau) \mapsto (v, ())], \psi^{\mathtt{rd}}[\pi \mapsto \sigma], \ldots, \sigma^{\mathtt{na}}[\ell \mapsto \tau] \rangle}{\langle \mathsf{E}[[\ell]_{\mathtt{na}} := v], \xi \rangle \Longrightarrow \langle \mathsf{E}[v], \xi' \rangle} \textsc{WriteNA}$$

$$\frac{\xi = \langle H, \psi^{\mathtt{rd}}, \ldots, \sigma^{\mathtt{na}} \rangle \quad \ldots \quad \tau = \mathsf{LastTS}(H, \ell) \quad \tau \equiv \sigma_{\mathtt{rd}}(\ell) \quad H(\ell, \tau) = (v, \sigma)}{\langle \mathsf{E}[[\ell]_{\mathtt{na}}], \xi \rangle \Longrightarrow \langle \mathsf{E}[v], \xi \rangle} \textsc{ReadNA}$$

$$\frac{\xi = \langle H, \psi^{\mathtt{rd}}, \ldots, \sigma^{\mathtt{na}} \rangle \quad \ldots \quad \sigma_{\mathtt{rd}}(\ell) \neq \mathsf{LastTS}(H, \ell)}{\langle \mathsf{E}[[\ell]_{\mathtt{na}} := v], \xi \rangle \Longrightarrow \langle \mathbf{stuck}, \xi \rangle} \textsc{WriteNA-stuck1}$$

$$\frac{\xi = \langle H, \psi^{\mathtt{rd}}, \ldots, \sigma^{\mathtt{na}} \rangle \quad \ldots \quad \sigma_{\mathtt{rd}}(\ell) \neq \mathsf{LastTS}(H, \ell)}{\langle \mathsf{E}[[\ell]_{\mathtt{na}}], \xi \rangle \Longrightarrow \langle \mathbf{stuck}, \xi \rangle} \textsc{ReadNA-stuck1}$$

$$\frac{\xi = \langle H, \psi^{\mathtt{rd}}, \ldots, \sigma^{\mathtt{na}} \rangle \quad \ldots \quad \sigma_{\mathtt{rd}}(\ell) < \sigma^{\mathtt{na}}(\ell)}{\langle \mathsf{E}[[\ell]_{\mathsf{RM}} := v], \xi \rangle \Longrightarrow \langle \mathbf{stuck}, \xi \rangle} \textsc{WriteNA-stuck2}$$

$$\frac{\xi = \langle H, \psi^{\mathtt{rd}}, \ldots, \sigma^{\mathtt{na}} \rangle \quad \ldots \quad \sigma_{\mathtt{rd}}(\ell) < \sigma^{\mathtt{na}}(\ell)}{\langle \mathsf{E}[[\ell]_{\mathsf{RM}}], \xi \rangle \Longrightarrow \langle \mathbf{stuck}, \xi \rangle} \textsc{ReadNA-stuck2}$$

Figure 25. Reduction rules for non-atomics.

$$\frac{\xi = \langle \ldots, \varphi, \gamma \rangle \quad \ldots \quad \mathsf{x} \text{ — fresh symbolic variable} \quad \alpha = \varphi(\pi) \quad \xi' = \langle \ldots, \varphi[\pi \mapsto \mathsf{append}(\alpha, \mathsf{E}\alpha, \mathit{read}\langle \mathsf{x}, \iota, \mathsf{RM} \rangle)], \gamma \rangle}{\langle \mathsf{E}[\mathsf{E}\alpha[[\iota]_{\mathsf{RM}}]], \xi \rangle \Longrightarrow \langle \mathsf{E}[\mathsf{E}\alpha[\mathsf{x}]], \xi' \rangle} \textsc{Read-Postpone}$$

$$\frac{\xi = \langle \ldots, \varphi, \gamma \rangle \quad \ldots \quad \mathsf{x} \text{ is fresh symbolic variable} \quad \alpha = \varphi(\pi) \quad \xi' = \langle \ldots, \varphi[\pi \mapsto \mathsf{append}(\alpha, \mathsf{E}\alpha, \mathit{write}\langle \mathsf{x}, \iota, \mathsf{WM}, \mathsf{e} \rangle)], \gamma \rangle}{\langle \mathsf{E}[\mathsf{E}\alpha[[\iota]_{\mathsf{WM}} := \mathsf{e}]], \xi \rangle \Longrightarrow \langle \mathsf{E}[\mathsf{E}\alpha[\mathsf{x}]], \xi' \rangle} \textsc{Write-Postpone}$$

$$\frac{\xi = \langle \ldots, \varphi, \gamma \rangle \quad \ldots \quad \mathsf{e} \text{ can't be substituted immediately} \quad \mathsf{x} \text{ — fresh symbolic variable} \quad \alpha = \varphi(\pi) \quad \xi' = \langle \ldots, \varphi[\pi \mapsto \mathsf{append}(\alpha, \mathsf{E}\alpha, \mathit{let}\langle \mathsf{x}, \mathsf{e} \rangle)], \gamma \rangle}{\langle \mathsf{E}[\mathsf{E}\alpha[\mathsf{x}' = \mathsf{e}; \mathsf{s}]], \xi \rangle \Longrightarrow \langle \mathsf{E}[\mathsf{E}\alpha[\mathsf{x}'/\mathsf{x}]], \xi' \rangle} \textsc{Let-Postpone}$$

$$\frac{\xi = \langle H, \psi^{\mathtt{rd}}, \ldots, \varphi, \gamma \rangle \quad \ldots \quad \mathit{read}\langle \mathsf{x}, \ell, \mathsf{RM} \rangle \text{ is inside } \varphi(\pi) \quad \text{and there is no conflicting operation before} \quad \varphi' = \mathrm{remove}(\varphi, \pi, \mathit{read}\langle \mathsf{x}, \ell, \mathsf{RM} \rangle) \quad \gamma' = \gamma \setminus \{\mathsf{x}\} \quad \xi' = \langle \ldots, \varphi'[v/\mathsf{x}], \gamma' \rangle}{\langle \mathsf{s}, \xi \rangle \Longrightarrow \langle \mathsf{s}[v/\mathsf{x}], \xi' \rangle} \textsc{Read-Resolve}$$

Figure 26. Rules for work with postponing of operations.

C. RCU testing

#	r11	r12	r21	r22	**Stuck**	**Runtime** (sec)
1	0	111	111	1101	✓	25.2
2	0	1	111	1101		21.4
3	0	0	0	0		12.9
4	0	1101	11	11		25.4
5	0	1101	0	0		16.3
6	0	11	0	0	✓	17.5
7	0	0	0	1101		22.1
8	1	1	0	0		16.5
9	0	1101	1	1101		19.2
10	0	1101	111	1101		26.4
11	1	1101	1	1		23.8
12	0	0	111	1101	✓	20.5
13	11	1101	0	0		21.5
14	0	111	0	111	✓	21.5
15	0	0	11	1101		22.1
16	0	0	0	0		16.0
17	0	0	0	1101		18.1
18	1	1	0	0		22.2
19	1	1101	1	1		26.0
20	1	1101	0	0		20.1

Figure 27. Test results and runtimes for modified RCU.